

Synthetic Landmine Scene Development and Validation in DIRSIG

by

Erin D. Peterson

B.S. Clarkson University

(1998)

A thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Chester F. Carlson Center for Imaging Science
of the College of Science
Rochester Institute of Technology

August 2004

Signature of the Author _____ <<Original Signed>>

Accepted by _____ <<Original Signed>> 26 Aug 2004

Coordinator, M.S. Degree Program

Date

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 04 OCT 2004		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Synthetic Landmine Scene Development and Validation in DIRSIG				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rochester Institute of Technology				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 148	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

CHESTER F. CARLSON
CENTER FOR IMAGING SCIENCE
COLLEGE OF SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

M.S. DEGREE THESIS

The M.S. Degree Thesis of Erin D. Peterson
has been examined and approved by the
thesis committee as satisfactory for the
thesis requirement for the
Master of Science degree

<<Original Signed>>

Dr. John Schott, Thesis Advisor

<<Original Signed>>

Dr. Carl Salvaggio

<<Original Signed>>

Dr. David Messinger

25 Aug 2004

Date

THESIS RELEASE PERMISSION
ROCHESTER INSTITUTE OF TECHNOLOGY
COLLEGE OF SCIENCE
CHESTER F. CARLSON
CENTER FOR IMAGING SCIENCE

Synthetic Landmine Scene Development and Validation in DIRSIG

I, Erin D. Peterson, do hereby grant permission to the Wallace Memorial Library of R.I.T. to reproduce my thesis in whole or in part. Any reproduction will not be for commercial use or profit.

Signature: _____ <<Original Signed>>_____

Date: _____ 25 Aug 2004_____

Synthetic Landmine Scene Development and Validation in DIRSIG

by
Erin D. Peterson

Submitted to the
Chester F. Carlson Center for Imaging Science
in partial fulfillment of the requirements
for the Master of Science Degree
at the Rochester Institute of Technology

ABSTRACT

Detection and neutralization of surface-laid and buried landmines has been a slow and dangerous endeavor for military forces and humanitarian organizations throughout the world. In an effort to make the process faster and safer, scientists have begun to exploit the ever-evolving passive electro-optical realm of detectors, both from a broadband perspective and a multi or hyperspectral perspective. Carried with this exploitation is the development of mine detection algorithms that take advantage of spectral features exhibited by mine targets, only available in a multi or hyperspectral data set. Difficulty in algorithm development arises from a lack of robust data, which is needed to appropriately test the validity of an algorithm's results. This paper discusses the development of synthetic data using the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model. A synthetic landmine scene has been modeled representing data collected at an arid US Army test site by the University of Hawaii's Airborne Hyperspectral Imager (AHI). The synthetic data has been created and validated to represent the surrogate minefield thermally, spatially, spectrally, and temporally over the 7.9 to 11.5 micron region using 70 bands of data. Validation of the scene has been accomplished by direct comparison to the AHI truth data using qualitative band to band visual analysis, radiance curve comparison, Rank Order Correlation comparison, Principle Components dimensionality analysis, Gray Level Co-occurrence Matrix and Spectral Co-occurrence Matrix analysis, and an evaluation of the R(x) algorithm's performance. This paper discusses landmine detection phenomenology, describes the steps taken to build the scene, modeling methods utilized to overcome input parameter limitations, and compares the synthetic scene to truth data.

Disclaimer

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

Acknowledgements

I would like to take this opportunity to thank my thesis committee, Dr. John Schott, Dr. Carl Salvaggio, and Dr. David Messinger, for their support and direction over the last year. Thank you to all the members of the DIRS staff and my fellow students for all the time, effort, and help you've given me. A big thanks goes to the contingency of Air Force students who have become like family during the course of this program. Finally, thank you to Cindy Schultz for all the phenomenal administrative support. Cindy keeps the DIRS group running like a well-oiled machine.

Dedication

This work is dedicated to my wonderful wife, my charismatic son, and my precious daughter. Your unquestionable love and unwavering support made this research and thesis possible.

Table of Contents

1.0	Introduction.....	1
2.0	Objectives and Work Statement.....	4
3.0	Background and Literature Review	5
3.1	Radiation Propagation.....	5
3.2	DIRSIG Overview	8
3.2.1	Scene Geometry Submodel.....	8
3.2.2	Ray-Tracer Submodel	9
3.2.3	Thermal Submodel.....	10
3.2.4	Radiometry Submodel	11
3.2.5	Sensor Submodel	12
3.2.6	DIRSIG Submodel Summary	13
3.2.7	DIRSIG Mapping Routines.....	14
3.3	Landmine Introduction.....	16
3.3.1	Anti-Tank Mines	17
3.3.2	Anti-Personnel Mines	18
3.4	Mine Emplacement	19
3.4.1	Minefields	20
3.5	Landmine Signatures	22
3.5.1	Surface Landmine Signatures	22
3.5.2	Buried Landmine Signatures.....	27
3.5.3	Signature Modeling.....	32
3.6	Comparison Metrics.....	32
3.6.1	Rank Order Correlation.....	33
3.6.2	Gray Level Co-occurrence Matrix	34
3.6.3	Spectral Co-occurrence Matrix	37
3.6.4	R(x) Anomaly Detection Algorithm	39
4.0	Approach.....	42
4.1	Microscene Experimental Data Collection	42
4.1.1	Timeline of Events.....	49
4.1.2	Collected Data.....	51
4.2	US Army Arid Test Site Scene	53
4.2.1	Available Ground Truth.....	54
4.2.2	AHI Radiance Images	57
4.3	DIRSIG Scene Construction	59
4.3.1	Terrain and Basic Scene Development	59
4.3.2	Target Development and Scene Geometry	61
4.3.3	Weather Data and Atmospheric Characterization.....	65
4.3.4	AHI Sensor Parameters.....	68
4.3.5	Material Parameters	83
4.3.6	Emissivity Extraction and Incorporation	92
4.3.7	Texture	98
5.0	Results.....	101

5.1	Visual Image Comparison.....	101
5.2	Rank Order Correlation.....	106
5.3	Radiance Curve Comparison	109
5.4	Dimensionality Analysis.....	114
5.5	GLCM and SCM Analysis.....	116
5.6	R(x) Algorithm Performance	122
6.0	Conclusions and Recommendations	127
6.1	Cyclical Modeling.....	127
6.2	Scene Improvements.....	129
6.3	Fundamental LWIR Modeling Improvements.....	130
6.4	Summary	131
7.0	References.....	132

List of Figures

Figure 3-1: Solar photon paths (Schott, 1997).....	5
Figure 3-2: Self-emitted photon paths (Schott, 1997)	5
Figure 3-3: Relative scene geometry within DIRSIG (DIRSIG homepage, 2003).....	9
Figure 3-4: Ray tracing within DIRSIG (Schott et al., 2001)	9
Figure 3-5: Ray tracing for downwelled radiance (Schott et al., 2001).....	10
Figure 3-6: DIRSIG submodel interactions.....	13
Figure 3-7: Bump mapping methodology.....	15
Figure 3-8: Neighboring pixels in a bump map.....	15
Figure 3-9: Components of a generic mine (Army, 2002).....	16
Figure 3-10: US Army M15 AT mine	18
Figure 3-11: U.S. Army M6A2 AT mine	18
Figure 3-12: Pressure activated AP mine, example #1	19
Figure 3-13: Pressure activated AP mine, example #2.....	19
Figure 3-14: Employment of AT pressure mine (Army, 2002)	20
Figure 3-15: Proper mine burial techniques (Army, 2002).....	20
Figure 3-16: Standard rapidly emplaced row minefield (Army, 2002)	21
Figure 3-17: Polarization scene set-up	26
Figure 3-18: Imaged polarization scene after vegetation growth	26
Figure 3-19: Linear polarization image	27
Figure 3-20: Daytime IR image of buried mines in soil.....	29
Figure 3-21: Nighttime IR image of buried mines in soil.....	29
Figure 3-22: Spectral signatures of disturbed and undisturbed soil.....	31
Figure 3-23: Spectral signatures for disturbed and undisturbed soil after flooding.....	31
Figure 3-24: GLCM example image	35
Figure 3-25: General GLCM form for image with four gray levels	35
Figure 3-26: GLCM example matrices.....	36
Figure 3-27: SCM base window from band #1	38
Figure 3-28: SCM shift window from band #2	38
Figure 3-29: Resulting spectral co-occurrence matrix.....	38
Figure 4-1: Microscene area	42
Figure 4-2: Microscene experimental set-up	44
Figure 4-3: Raised scissor cart.....	45
Figure 4-4: Mine area in Microscene.....	46
Figure 4-5: SIMS test set, courtesy SIMS homepage	47
Figure 4-6: Mine area layout	47
Figure 4-7: Detailed mine area layout	48
Figure 4-8: WASP Visible at 1400hrs	51
Figure 4-9: WASP LWIR at 0200hrs	51
Figure 4-10: MISI VIS Image, 1000hrs.....	52
Figure 4-11: Collection area	53
Figure 4-12: Top Hat fiducial target.....	54
Figure 4-13: Photo capture locations.....	55
Figure 4-14: Metallic mine (L); Desert Bush (R)	56
Figure 4-15: Non-roll corrected AHI images, 11.0 μ m (top) and 9.2 μ m (bottom), (Winter, 2003)	59
Figure 4-16: Bump map gray level image	60
Figure 4-17: Metallic mine model.....	62
Figure 4-18: Desert Bush model.....	62
Figure 4-19: Scene target placement and coordinate system.....	63
Figure 4-20: Disturbed soil area in material map	64
Figure 4-21: Disturbed ground base material map	65

Figure 4-22: V/H Error, Side View	70
Figure 4-23: V/H Error, Top View	70
Figure 4-24: Illustration of the compression effect	70
Figure 4-25: 3x oversampling example	71
Figure 4-26: Inter-pixel blur illustration.....	73
Figure 4-27: AHI noise development flowchart - Step1.....	76
Figure 4-28: AHI noise development flowchart - Step 2.....	77
Figure 4-29: (A) E700 AHI derived noise covariance, (B) E700 synthetic noise covariance	78
Figure 4-30: (A) E1400 AHI derived noise covariance, (B) E1400 synthetic noise covariance	78
Figure 4-31: (A) N700 AHI derived noise covariance, (B) N700 synthetic noise covariance	78
Figure 4-32: (A) N1400 AHI derived noise covariance, (B) N1400 synthetic noise covariance	78
Figure 4-33: E700 Synthetic noise SNR.....	79
Figure 4-34: E1400 Synthetic noise SNR.....	79
Figure 4-35: N700 Synthetic noise SNR	79
Figure 4-36: N1400 Synthetic noise SNR	79
Figure 4-37: AHI truth SNR plot (Lucey and Winter, 1998)	79
Figure 4-38: AHI E700, Band 25	81
Figure 4-39: DIRSIG E700, Band 25	81
Figure 4-40: DIRSIG roll corrected E700, Band 25	83
Figure 4-41: (L) AHI temperature profile, E700; (R) Initial DIRSIG temperature profile, E700	86
Figure 4-42: (L) AHI temperature profile, E1400; (R) Initial DIRSIG temperature profile, E1400	87
Figure 4-43: (L) AHI temperature profile, N700; (R) Initial DIRSIG temperature profile, N700	87
Figure 4-44: (L) AHI temperature profile, N1400; (R) Initial DIRSIG temperature profile, N1400	87
Figure 4-45: Material map with compact disturbed ground areas	88
Figure 4-46: Undisturbed ground variant mapping image.....	89
Figure 4-47: (L) AHI temperature profile, E700; (R) Final DIRSIG temperature profile, E700	90
Figure 4-48: (L) AHI temperature profile, E1400; (R) Final DIRSIG temperature profile, E1400.....	90
Figure 4-49: (L) AHI temperature profile, N700; (R) Final DIRSIG temperature profile, N700.....	91
Figure 4-50: (L) AHI temperature profile, N1400; (R) Final DIRSIG temperature profile, N1400.....	91
Figure 4-51: Derived atmospheric transmission and upwelled radiance	94
Figure 4-52: (L) Blackbody curve fit for undisturbed soil, (R) Resulting emissivity curve.....	96
Figure 4-53: Derived emissivity of disturbed and undisturbed soil.....	97
Figure 4-54: Emissivity of undisturbed and disturbed soil (Winter et al., 1996).....	97
Figure 4-55: Tileable texture image	99
Figure 5-1: Noon at 700 feet, band 25, (L) AHI, (R) DIRSIG	102
Figure 5-2: Noon at 700 feet, band 65, (L) AHI, (R) DIRSIG	102
Figure 5-3: Noon at 1400 feet, band 25, (L) AHI, (R) DIRSIG	102
Figure 5-4: Noon at 1400 feet, band 65, (L) AHI, (R) DIRSIG	102
Figure 5-5: Evening at 700 feet, band 25, (L) AHI, (R) DIRSIG.....	103
Figure 5-6: Evening at 700 feet, band 65, (L) AHI, (R) DIRSIG.....	103
Figure 5-7: Evening at 1400 feet, band 25, (L) AHI, (R) DIRSIG	103
Figure 5-8: Evening at 1400 feet, band 65, (L) AHI, (R) DIRSIG.....	103
Figure 5-9: Image data histograms from band 25, midday at 700 feet, (L) AHI, (R) DIRSIG.....	104
Figure 5-10: Image data histograms from band 25, evening at 1400 feet, (L) AHI, (R) DIRSIG.....	104
Figure 5-11: Radiance curve comparison of evening data at 700 feet, (L) AHI, (R) DIRSIG	110
Figure 5-12: Radiance curve comparison of evening data at 1400 feet, (L) AHI, (R) DIRSIG	110
Figure 5-13: Radiance curve comparison of midday data at 700 feet, (L) AHI, (R) DIRSIG.....	112
Figure 5-14: Radiance curve comparison of midday data at 1400 feet, (L) AHI, (R) DIRSIG.....	112
Figure 5-15: Cumulative variance plot - E700	115
Figure 5-16: Cumulative variance plot - E1400	115
Figure 5-17: Cumulative variance plot - N700.....	115
Figure 5-18: Cumulative variance plot - N1400.....	115
Figure 5-19: E700, band 26 GLCM subset, (L) AHI, (R) DIRSIG	117

Figure 5-20: E1400, band 26 GLCM subset, (L) AHI, (R) DIRSIG	117
Figure 5-21: N700, band 26 GLCM subset, (L) AHI, (R) DIRSIG.....	117
Figure 5-22: N1400, band 26 GLCM subset, (L) AHI, (R) DIRSIG.....	117
Figure 5-23: E700, band 60 GLCM subset, (L) AHI, (R) DIRSIG	117
Figure 5-24: E1400, band 60 GLCM subset, (L) AHI, (R) DIRSIG	117
Figure 5-25: N700, band 60 GLCM subset, (L) AHI, (R) DIRSIG.....	117
Figure 5-26: N1400, band 60 GLCM subset, (L) AHI, (R) DIRSIG.....	117
Figure 5-27: R(x) result images for (L) AHI data and (R) DIRSIG data, midday at 700 feet.....	123
Figure 5-28: R(x) result images for (L) AHI data and (R) DIRSIG data, midday at 1400 feet	123
Figure 5-29: R(x) result images for (L) AHI data and (R) DIRSIG data, evening at 700 feet	124
Figure 5-30: R(x) result images for (L) AHI data and (R) DIRSIG data, evening at 1400 feet	124
Figure 5-31: R(x) ROC curve, evening at 700 feet.....	126
Figure 5-32: R(x) ROC curve, evening at 1400 feet.....	126
Figure 5-33: R(x) ROC curve, midday at 700 feet	126
Figure 5-34: R(x) ROC curve, midday at 1400 feet	126
Figure 6-1: Pictorial description of the cyclical scene building process used for the landmine scene	128

List of Tables

Table 3-1: Input parameters to THERM.....	10
Table 4-1: Portion of WAAMD weather listing	57
Table 4-2: Scene object dimensions	61
Table 4-3: Portion of the DIRSIG weather file.....	66
Table 4-4: Evening 700 ft convolution kernel	74
Table 4-5: Midday 700 ft convolution kernel	74
Table 4-6: Evening 1400 ft convolution kernel	74
Table 4-7: Midday 1400 ft convolution kernel.....	74
Table 4-8: Material map descriptions	90
Table 5-1: Number of pixels per material used in ROC evaluation.....	107
Table 5-2: Overall scene rank order correlation statistics, comparing DIRSIG to AHI imagery	108
Table 5-3: Individual material rank order correlation values, comparing DIRSIG to AHI imagery	109
Table 5-4: A comparison of data variance between AHI and DIRSIG across all sets of data.....	114
Table 5-5: GLCM contrast statistics from band 26, AHI data.....	118
Table 5-6: GLCM contrast statistics from band 26, DIRSIG data.....	118
Table 5-7: GLCM contrast statistics from band 60, AHI data.....	118
Table 5-8: GLCM contrast statistics from band 60, DIRSIG data.....	118
Table 5-9: GLCM contrast statistics with large window from band 26, AHI data	119
Table 5-10: GLCM contrast statistics with large window from band 26, DIRSIG data.....	119
Table 5-11: GLCM contrast statistics with large window from band 60, AHI data	120
Table 5-12: GLCM contrast statistics with large window from band 60, DIRSIG data.....	120
Table 5-13: SCM contrast statistics from AHI data	121
Table 5-14: SCM contrast statistics from DIRSIG data	121
Table 5-15: SCM contrast statistics with large window from AHI data.....	121
Table 5-16: SCM contrast statistics with large window from DIRSIG data.....	121

1.0 Introduction

Since World War II, the ability of an army to detect a minefield that lay in wait has literally meant the difference between life and death. Not only does this apply to advancing forces during a conflict, but also to the residents of the region after the conflict has ended. This introduces the concept of humanitarian demining, a process to ensure innocent civilians are not seriously injured as a result of forgotten and undetectable minefields. As technology has advanced over the decades, the ability to collect and exploit a wider range of data pertaining to landmines and minefields has advanced as well. This has given scientists new abilities to attempt detection of mines and minefields. With a new arsenal of information at the disposal of scientists, the development of new methods to detect mines has followed. From simple metal detectors to the use of thermal imaging or ground penetrating radar, advances have taken place in the countermining field at a quick pace. The challenge at the forefront of research is testing these various detection techniques to quantify how well each performs or potentially, how well some perform in tandem.

The introduction of novel detection techniques into the countermining community has created a need for test data that provides a safe environment for researchers, without the loss of critical information about a mine or a minefield's "signature". By limiting this discussion of the countermining field to passive electro-optical detection techniques, Synthetic Image Generation (SIG) may prove to be the solution to the problem at hand.

Researchers developing anomaly or target detection algorithms for use on hyperspectral data sets or broadband images, cannot afford to undertake huge experimental efforts to produce data spanning a multitude of imaging conditions for evaluation of algorithm performance. It is simply too cost prohibitive and work intensive. Therefore, the need for accurate, reproducible images or sets of data is paramount. SIG can accomplish this task. Specifically, the Digital Imaging and Remote Sensing (DIRS) group's Image Generation (DIRSIG) model can be used for this purpose. DIRSIG is a first-principles based physics model that produces high fidelity radiance images of a synthetic scene. Using DIRSIG to build synthetic images offers the

algorithm developer total control and flexibility over the data produced. Therefore, if DIRSIG accurately models the physical interactions between objects within the scene, algorithm developers can use the flexibility of the scene generation process to evaluate algorithm performance in a seemingly endless set of scenarios.

The purpose of this project is to use DIRSIG to produce high-resolution images of a scene that contains surface-laid and buried landmine signatures as well as applicable background objects and clutter. The scene can then be used by algorithm developers as a “training ground” to test algorithm performance. Correctly modeling the physics behind object interactions within the scene is of critical importance. If the underlying principles are correct, a scene with accurate minefield signatures will be created, and mine detection or anomaly detection algorithms will perform equally as well on synthetic and real-world data. Hopefully, by using robust data to train an algorithm or refine an algorithm, a more robust algorithm will emerge.

The work accomplished for this project stems from a U.S. Army, Multi-University Research Initiative (MURI) sponsored by the Army Research Organization (ARO). Together with four other universities, RIT hopes to further the understanding of the science behind today’s ever-growing landmine detection problem. The lead organization for the MURI project is Georgia Institute of Technology and their piece of the overarching puzzle is phenomenology exploitation. In addition to Georgia Tech, the University of Maryland is responsible for Automatic Target Recognition (ATR) algorithm development, the University of Florida is responsible for data fusion, the University of Hawaii is responsible for data collection and processing, and finally RIT is responsible for synthetic scene generation. RIT has taken data collected from the University of Hawaii’s AHI sensor and reproduced a representative landmine scene. The scene will be used by the University of Maryland to robustly test their mine detection algorithms.

The approach described in this effort deals not with the specifics of algorithm development, but the generation of accurate synthetic data. A high-resolution scene has been developed and validated by comparing minefield and background attributes of the

synthetic scene to attributes of a known data set. Quantification of the differences between the synthetic and the truth has also been accomplished. The goal of this work was to provide a validated scene that accurately represents landmine signatures to an ATR algorithm. The following sections will detail the background and approach that was used for generation and validation of the scene.

2.0 Objectives and Work Statement

The objectives of this research were to:

1. Research and report about standard deployment of surface and buried mines; research and report about surface and buried mine signatures utilized by ATR algorithms for mine detection.
2. Design and perform an experiment to collect MWIR and LWIR landmine signature data utilizing MISI and WASP imaging systems on RIT's campus.
3. Create/build a high-resolution DIRSIG scene based on data collected at an arid US Army test site for the Army MURI project.
4. Adequately represent surface and buried landmine signatures within the scene.
5. Render the scene at resolution and viewing angles comparable to truth data under approximately four different viewing conditions, e.g. time of day, sensor viewing angle, or sensor elevation.
6. Validate landmine signatures in DIRSIG rendered scenes using quantifiable metrics, using the AHI data as truth. Metrics to be used are Rank Order Correlation, Gray Level Co-occurrence Matrix and Spectral Co-occurrence Matrix evaluation with emphasis on the contrast metric, and evaluation of R(x) algorithm's performance.
7. Based on initial evaluation, propose and implement improvements to the synthetic scene that will enhance the accuracy of landmine signatures.

3.0 Background and Literature Review

This section introduces the reader to the background concepts and theory used in this project. Included will be a review of radiation transfer theory and an introduction to DIRSIG and its components. A general introduction to landmines and the theory behind landmine usage in wartime environments will follow. Then, a discussion of landmine signatures will be presented followed by a discussion of the metrics used to validate the scene.

3.1 Radiation Propagation

To understand how any SIG model renders radiance images, more specifically DIRSIG, a review of the basics behind radiance propagation to the imaging sensor is required. At the outset, radiation that reaches the sensor may be split into two types, solar produced radiance and self-emitted radiance. This treatment concentrates on passive imaging, so a discussion pertaining to active illumination of a target is not presented. Figure 3-1 shows five possible paths that solar photons may travel enroute to the sensor. The notation throughout this discussion will be consistent with the notation presented by Schott (1997).

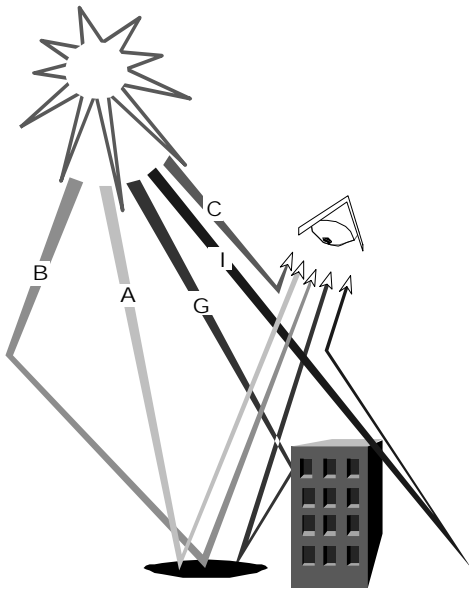


Figure 3-1: Solar photon paths (Schott, 1997)

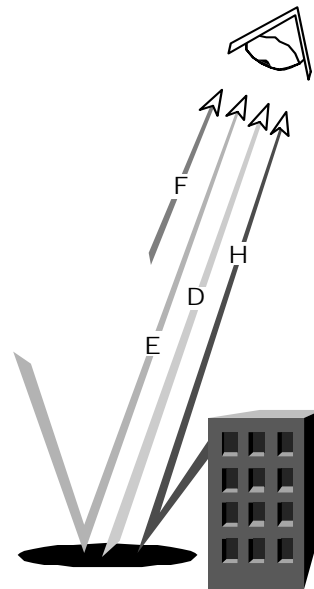


Figure 3-2: Self-emitted photon paths (Schott, 1997)

Mathematically, total solar produced radiance reaching the sensor (L_{solar}) can be described in the following fashion,

$$L_{\text{solar}} = L_A + L_B + L_C + L_G \quad (3-1)$$

where,

- L_A solar radiance passed through the atmosphere, reflected off the target, passed through the atmosphere to the sensor
- L_B solar radiance scattered by the atmosphere onto the target, reflected off the target, passed through the atmosphere to the sensor
- L_C solar radiance scattered by the atmosphere directly to the sensor
- L_G solar radiance passed through the atmosphere, reflected off a background object onto the target, reflected off the target, passed through the atmosphere to the sensor

I-type photons have been attenuated by the atmosphere, reflected off a background object, and then scattered by the atmosphere to the sensor. According to Schott (1997), I-type photons can be grouped with C-type photons if the average albedo of the scene is slowly varying, which is common. In the equation presented, this assumption had been made hence the absence of the L_I term.

Self-emitted photon paths are shown in Figure 3-2. These paths describe photon travel while imaging in the thermal region of the spectrum (MWIR, LWIR).

Mathematically, self-emitted thermal radiance (L_{thermal}) is represented by,

$$L_{\text{thermal}} = L_D + L_E + L_F + L_H \quad (3-2)$$

where,

- L_D self-emitted radiance from the target, through the atmosphere to the sensor
- L_E self-emitted radiance from the atmosphere, reflected off the target, through the atmosphere to the sensor
- L_F self-emitted radiance from the atmosphere directly to the sensor
- L_H self-emitted radiance from a background object, reflected off the target, through the atmosphere to the sensor

Combining solar radiance reaching the sensor with self-emitted radiance reaching the sensor, the wavelength dependent “Big Equation” is formed, describing the total radiance reaching the sensor.

$$L_{\lambda} = L_A + L_D + L_B + L_E + L_G + L_H + L_C + L_F \quad (3-3)$$

Substituting in expressions as given in Schott (1997) for each component results in the full mathematical version of the Big Equation,

$$L_{\lambda} = \left\{ E_{s\lambda} \cos \sigma \tau_1(\lambda) \frac{r(\lambda)}{\pi} + \varepsilon(\lambda) L_{T\lambda} + F [E_{ds\lambda} + E_{de\lambda}] \frac{r_d(\lambda)}{\pi} + (1-F) [L_{bs\lambda} + L_{be\lambda}] r_d(\lambda) \right\} \tau_2(\lambda) + L_{us\lambda} + L_{ue\lambda} \quad (3-4)$$

where,

$E_{s\lambda}$	exoatmospheric irradiance
$\cos \sigma$	angle from target normal to the sun
$\tau_1(\lambda)$	transmission through atmosphere along sun-target path
$r(\lambda)$	target reflectance
$\varepsilon(\lambda)$	target emissivity
$L_{T\lambda}$	self-emitted radiance from target at temperature T
F	fraction of hemisphere above target that is sky (shape factor)
1-F	fraction of hemisphere above target that is background
$E_{ds\lambda}$	downwelled solar irradiance
$E_{de\lambda}$	downwelled self-emitted radiance from atmosphere
$r_d(\lambda)$	target diffuse reflectance
$\tau_2(\lambda)$	transmission through atmosphere along target-sensor path
$L_{bs\lambda}$	background reflected solar radiance onto target
$L_{be\lambda}$	background self-emitted radiance onto target
$L_{us\lambda}$	upwelled solar irradiance
$L_{ue\lambda}$	upwelled self-emitted radiance

The wavelength dependency of sensor reaching radiance shows the solution to the big equation will vary depending on the region of the spectrum being imaged. If

concerned only with visible or near infrared wavelengths, the self-emitted portions of the Big Equation can be neglected with minimal error. Complementary to this, if imaging in the long wave infrared, the solar terms may be neglected with minimal error.

3.2 DIRSIG Overview

Scene modeling and validation has been accomplished using the Digital Imaging and Remote Sensing (DIRS) group's Image Generation (DIRSIG) model. DIRSIG is a first-principles based model, which mathematically represents the entire imaging chain to produce radiance images in the 0.3 to 20 micron region of the electromagnetic spectrum. The model is modular in design, such that the DIRSIG executable program combines the results of individual sub-models, each responsible for a link in the image chain. Six basic submodels are utilized by the DIRSIG executable namely the scene geometry submodel, ray-tracer submodel, thermal submodel, radiometry submodel, sensor submodel, and the plume generation submodel. As plume generation is not applicable to this work, the plume generation submodel will not be discussed. A brief description of each submodel is presented.

3.2.1 Scene Geometry Submodel

The scene geometry submodel is a means to provide the three-dimensional description of the synthetic scene. Every object within the scene is generated using an enhanced CAD environment, namely Rhinoceros (a drawing program similar to AutoCAD) or purchased from a commercial drawing company. Once objects are drawn to scale, they are facetized, i.e. segmented into groups of two-dimensional polygons. Detailed objects may contain tens of thousands of individual facets. Each facet is assigned thermodynamic and optical properties coinciding with the material of the facet at hand, using a locally developed program called Bulldozer. In addition to facetized objects, scene geometry also accounts for relative positioning of all objects within the scene to include the imaging sensor. Once the scene is created, it can be imaged. Figure 3-3 shows relative positioning of individual scene facets in reference to the focal plane of the detector used to image the scene.

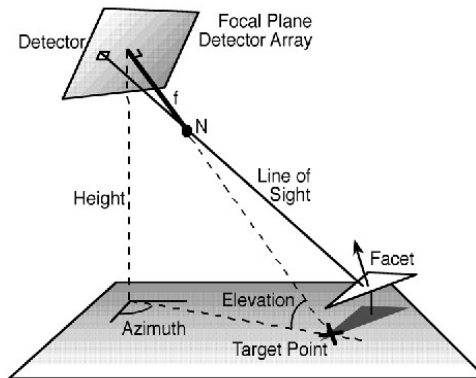


Figure 3-3: Relative scene geometry within DIRSIG (DIRSIG homepage, 2003)

3.2.2 Ray-Tracer Submodel

The scene's geometric data is sampled using standard ray-tracing methodology, taking into account data about the sensor's focal plane contained in the sensor submodel. In its most basic form, a ray is cast from each pixel in the sensor's focal plane into the scene. As the ray encounters facets in the scene, an interaction list is produced and recorded. Rays are then cast from the encountered facet to the sun in order to determine solar shadowing at the current time and solar loading information for up to 24 previous hours. This allows for accurate temperature predictions to be calculated. Figure 3-4 shows the ray tracing methodology used within DIRSIG as well as ray tracing to determine solar loading for the encountered facet.

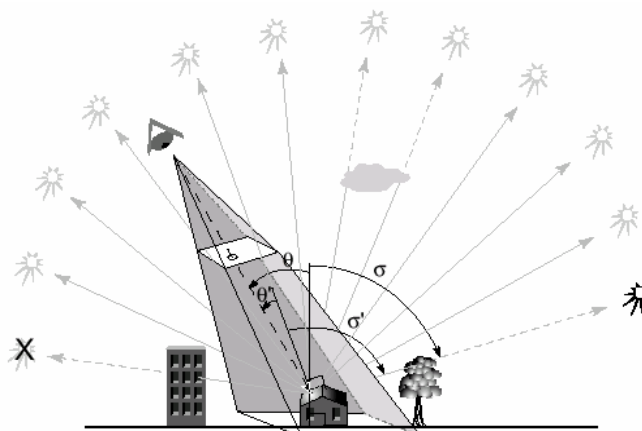


Figure 3-4: Ray tracing within DIRSIG (Schott et al., 2001)

Additional rays are cast from the facet into the hemisphere above the target for characterization of the downwelled radiance onto the target with respect to the material's

bi-directional reflectance distribution factor (BRDF). Figure 3-5 shows additional rays being cast from an encountered facet to determine the downwelled radiance field.

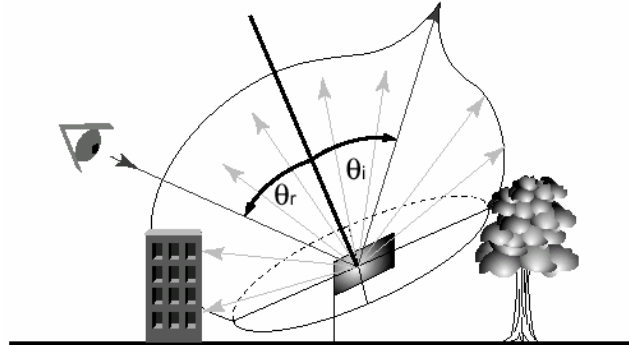


Figure 3-5: Ray tracing for downwelled radiance (Schott et al., 2001)

The reader is referred to Brown et al. (1997) for further detail on BRDF calculations in DIRSIG.

3.2.3 Thermal Submodel

The main purpose of the thermal submodel is to calculate a diurnal temperature prediction for each facet. To do this, DIRSIG employs THERM; a first-principles based model that evaluates heat transfer between a facet and its surround (DCS Corporation, 1991). A detailed set of parameters is given to the model, which include geographic parameters, material parameters, and meteorological parameters (listed in Table 3-1, adapted from Kraska, 1996 and Joseph, 1998).

<u>Geographic Parameters</u>	<u>Material Parameters</u>
Latitude	Heat Capacity
Longitude	Thermal Conductivity
Date	Thickness
Time Difference from GMT	Visible and Thermal Emissivity
Time Interval	Self Generated Power
Sunrise	Slope and Azimuthal Angles
Sunset	Shape Factor
	Exposed Area
<u>Meteorological Parameters</u>	
Direct and Diffuse Insolation	Sky Exposure
Air Temperature	Cloud Type
Air Pressure	Precipitation Type/Rate/Temp
Relative Humidity	Wind Speed
Dew Point	

Table 3-1: Input parameters to THERM

THERM evaluates the solar shadow history for each facet over the previous 24 hours in 15-minute intervals, to determine an accurate temperature prediction. Each facet is treated independently, which leads to a model limitation. Conduction between facets is not calculated. THERM has been adapted to include indirect, one-dimensional diffusion of internal heat sources, which does somewhat alleviate the limitation. In the case where extremely detailed temperature information of an object is needed, DIRSIG has the capability to incorporate an offline thermal calculation of scene objects. For example, a more sophisticated finite-element thermal model may be used to compute instantaneous temperature predictions for an object's facets. These individual facet temperatures would be inserted directly into DIRSIG. The drawback to thermal modeling this way is the advantage of background interaction with the object is negated. Once THERM completes temperature predictions for each facet, all information is passed to the radiometry submodel.

3.2.4 Radiometry Submodel

The radiometry submodel is the workhorse under DIRSIG's hood. The detailed calculations for determining radiance reaching the sensor, as described in Section 3.1, are performed by this submodel. Utilizing the information passed from the ray-tracer and the thermal model along with MODTRAN (Berk et al., 1999) or FASCODE (Smith et al., 1978), radiance reaching the sensor is computed for each pixel in the scene. MODTRAN is a U.S. Air Force developed program that computes atmospheric radiation propagation based on user inputs. FASCODE is a high resolution equivalent to MODTRAN. MODTRAN computes transmission as a function of view angle, upwelled and downwelled spectral radiance as a function of view angle, slant path, and range for any given sensor geometry (Schott et al., 2001). "Canned" atmosphere descriptions exist in MODTRAN, such as "tropical" or "mid-latitude summer" that can be rapidly integrated into a scene. If a more detailed atmospheric description is needed, user supplied radiosonde data can be incorporated.

The radiometry submodel is capable of accounting for transmissive facets as well as opaque facets. For transmissive facets, the ray-tracer develops a path through the facet

and then MODTRAN computes the radiance along the path. For opaque objects, BRDF calculations are used to determine specularity or diffusivity in the direction of the sensor. Background interactions are also calculated within this submodel, whether they are specular or diffuse contributions. A facet's self-emitted radiance is computed by solving the Planck equation (3-5) using temperature predictions output from the thermal submodel.

$$L(\lambda) = \frac{2hc^2}{\lambda^5 (e^{\frac{hc}{\lambda kT}} - 1)} \quad (3-5)$$

where,

- $L(\lambda)$ – spectral radiance
- T – absolute temperature of the object
- h – Plank's constant
- k – Boltzmann gas constant
- c – speed of light

3.2.5 Sensor Submodel

The final submodel describes the imaging sensor. Radiance at the sensor, as determined by the radiometry submodel, is converted to digital count values, corresponding to the sensor's spectral response function. Detailed geometric descriptions of sensors (e.g. line scanners, pushbroom scanners, or framing arrays) can be modeled. The radiance observed over a specified bandpass is accounted for by combining the radiance reaching the sensor with the sensor's spectral response function in the following manner (Joseph, 1998).

$$L_w(\theta) = \sum_{\lambda_{\min}}^{\lambda_{\max}} L(\theta, \lambda) R(\lambda) \Delta\lambda \quad (3-6)$$

where,

- $L_w(\theta)$ – radiance at zenith angle θ over the bandpass of interest
- $L(\theta, \lambda)$ – radiance at the sensor, at wavelength λ
- $R(\lambda)$ – normalized spectral response function of the sensor at wavelength λ

λ_{\max} & λ_{\min} – defines the bandpass of interest

$L_w(\theta)$ is then converted to a digital count by using a simple linear relationship.

$$DC = G \cdot L_w(\theta) + B \quad (3-7)$$

where,

DC – digital count

G – gain of sensor

B – bias of sensor

Additionally, more complex sensor interactions can be handled as well, such as geometric effects specific to sensor type, platform motion and others. The reader is encouraged to read Salacain (1995) for a detailed treatment of sensor characteristics in DIRSIG.

3.2.6 DIRSIG Submodel Summary

In short, the interaction between submodels is adequately represented in the following flowchart.

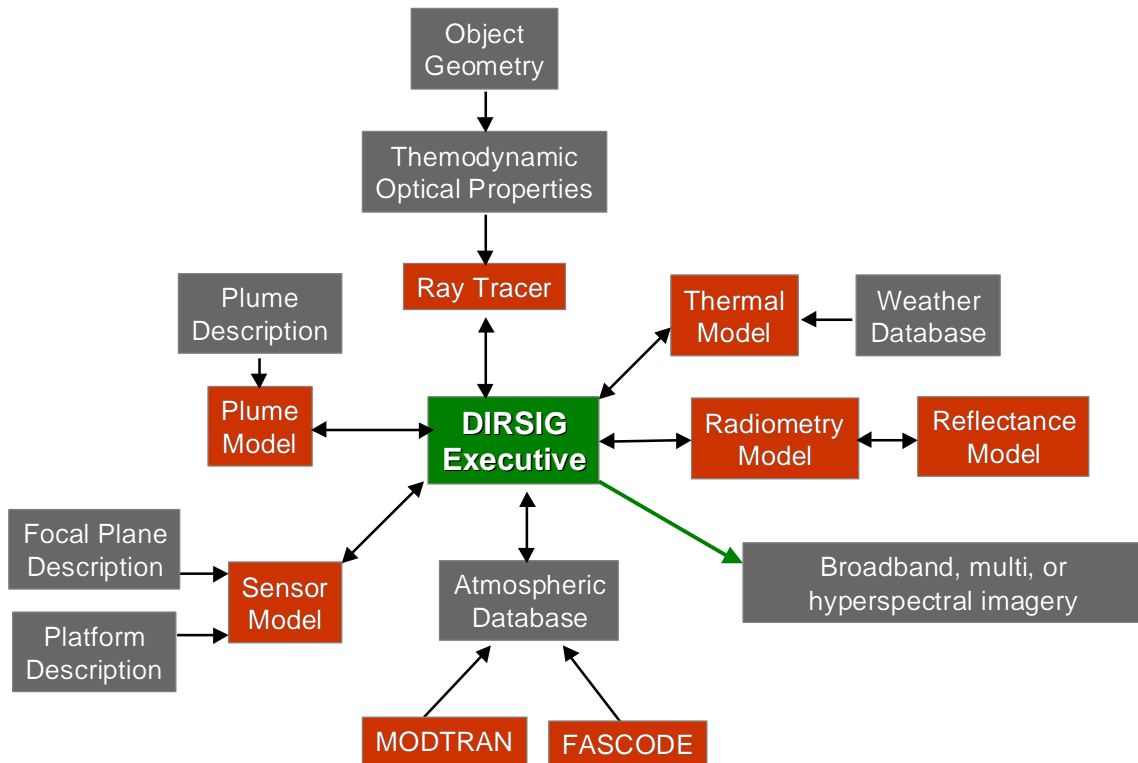


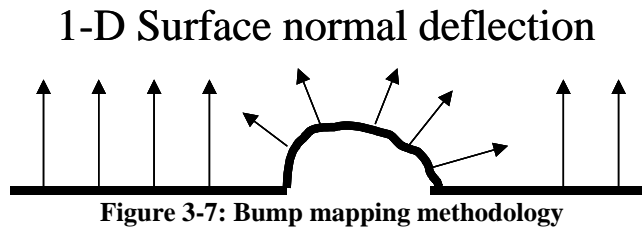
Figure 3-6: DIRSIG submodel interactions

3.2.7 DIRSIG Mapping Routines

In addition to DIRSIG's main submodels, there are built in mapping routines that give scene builders greater ability to project sub-facet variation into their work. This variation may include material variation, emissivity variation within a single material, temperature variation, material mixture variation for individual pixels, reflectance variation, or radiance variation (Brown and Schott, 2000). A "map", regardless of purpose, is typically a gray-level image that spatially represents some distribution of data within the scene. Each map has associated with it an insert point and a GSD value. The insert point tells DIRSIG where to orient the (0,0) point of the map image in relation to the scene's geometry. The GSD (ground sample distance) value tells DIRSIG how much area an individual mapping image pixel should cover in DIRSIG scene units. Each map will therefore have a ground resolution that can be input by the scene creator. This allows the scene creator to insert small details into a scene without having to create hundreds of thousands of tiny facets in order to increase the resolution of the scene. The use of mapping routines alleviates the headache of small facet generation while enhancing DIRSIG's ability to model small detail. The reader is encouraged to reference the DIRSIG user's manual for a complete treatment of mapping routines in DIRSIG.

Two mapping routines used in this work are material mapping and emissivity (texture) mapping. The use of both routines is extensively documented in the DIRSIG user's manual and will be discussed at length in following sections. However, in addition to material and texture mapping, a new mapping routine was used that is not well documented. Bump mapping allows a scene creator to insert two-dimensional surface variation into a given material. A gray-level mapping image is used where the difference in digital counts between neighboring mapping image pixels is used to alter the direction of a material's surface normal vector. For example, if a wooden picnic table were built using a CAD program, it would most likely have a perfectly flat and smooth surface. In reality, a wooden picnic table has many surface grooves and ridges that would not be modeled unless extraordinary time was taken to create them. A bump map is used to create the appearance of the bumps and ridges without having to build individual facets.

The map image would be applied to the table's surface. The amount of change in gray level value between two of the image's pixels would correspond to a deflection of the table's surface normal at that spatial location by a calculated angular amount. This concept is shown in Figure 3-7.



As with other mapping routines, the resolution of the map is user determined though the associated GSD value. Bump mapping also gives the user the ability to determine the range of angular deflection of the surface normal through an assigned scale value. This value ranges from 0 to 1. To illustrate the relationship between scale value and angular deflection, consider two neighboring pixels in a bump map image, as depicted in Figure 3-8.

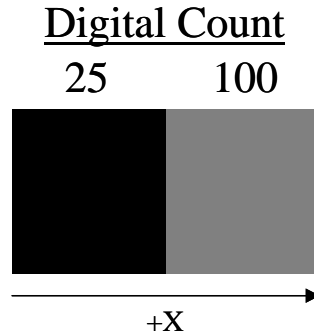


Figure 3-8: Neighboring pixels in a bump map

In a gray level image there are 256 possible digital count values. Also, a surface normal vector can be deflected by at most 90° ; therefore a change in digital count (Δx in this representation) by 255 corresponds to a deflection of 90° . The scale value associated with the map attenuates the deflection angle, limiting the maximum range that a normal vector can be deflected. The routine determines the change in DC value through Equation (3-8),

$$\Delta x = \text{image}[x + 1] - \text{image}[x] \quad (3-8)$$

A positive Δx indicates a normal deflection to the left, a negative Δx to the right. In the example from Figure 3-8, $\Delta x = 100 - 25 = 75$, which corresponds to a maximum angular deflection of $\sim 26.37^\circ$ to the left. The calculation works exactly the same in the Y direction. The use of this routine allows roughened surfaces to be introduced while bypassing time-consuming detailed object creation. The detail of the surface structure is only limited by the resolution of the mapping image itself.

3.3 Landmine Introduction

Now that the modeling platform has been introduced, the focus shifts to the data that is represented within the model. The first step to creating synthetic data for use by an Automatic Target Recognition (ATR) algorithm is to understand what target features the algorithm will key on. In this case, landmines and minefields are the targets, so naturally a minimum of a cursory introduction to landmines must be presented.

According to the U. S. Department of the Army, a landmine is defined as, “an explosive device that is designed to destroy or damage equipment or personnel.” (Army, 2002) A mine can be water-based or land-based and this discussion will be limited to land-based mines, hence landmines. A generic mine is a very simple mechanism comprised of five main components, a firing mechanism, detonator, booster charge, main charge and casing. A pictorial description is presented in the following figure.

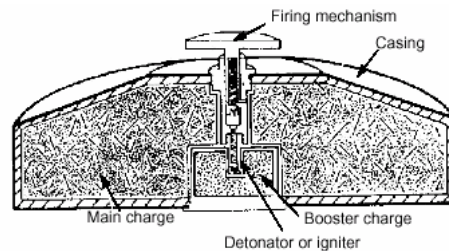


Figure 3-9: Components of a generic mine (Army, 2002)

The firing mechanism is the means for the mine to detonate. It may take many forms such as pressure activated, pull activated (tripping a trip-wire), time delay, vibration sensitive, or even infrared sensor activated, to name a few. Predominantly, pressure activated and pull activated mines are the most common due to their simplicity and

inexpensive production costs. Once the firing mechanism has been activated, the detonator simply ignites the booster charge, which provides the energy necessary to detonate the main charge of the mine. Since mines are designed to be compact weapons, the main charge usually comprises most of the mine body. The casing is simply the mine's outer shell. Originally, mine casings were made of metal or steel, which led to early mine detection with a simple metal detector that one might purchase at a hardware store. As sturdy plastics have been developed, mines have become increasingly less metallic. Not only has the introduction of plastic casings made mines easier and cheaper to produce, they have also made mine detection much more difficult from a traditional mine detection standpoint.

3.3.1 Anti-Tank Mines

Landmines are generally classified into two groups, anti-personnel (AP) or anti-tank (AT). AT mines are typically larger in size, due to the increased blast effects needed to incapacitate an armored vehicle. Additionally, AT pressure activated mines are typically designed so that a person stepping on it will not cause detonation. According to Bonsor (2001), the range of required applied weight to detonate an AT mine ranges from approximately 350 to 750 pounds. Due to weight requirements, these mines are typically emplaced on roads or in areas where vehicular traffic may occur. Actual sizes of AT mines range widely from roughly 20cm to over a meter in diameter. Shapes can vary from circular to quite irregularly shaped. Figure 3-10 and Figure 3-11 show U.S. Army, anti-tank landmines, courtesy of the Department of Defense and Norwegian People's Aid (NPA), as examples. These two mines are roughly 340 cm in diameter with the M15 at 125 cm in height and the M6A2 at 80 cm in height.



Figure 3-10: US Army M15 AT mine



Figure 3-11: U.S. Army M6A2 AT mine

3.3.2 Anti-Personnel Mines

Most AP mines fall into three general categories, blast, bounding, and fragmentation. Blast mines are the most common and detonate following activation of the firing mechanism. Effects of the explosive blast are the main destructive force, however, secondary damage is caused by the fragmentation of the mine casing.

Bounding mines are different in that upon activation of the firing mechanism, an initial charge is fired, hurling the mine up to a meter in the air. The main charge is then detonated, releasing multiple fragments into the immediate area. These fragments combined with blast effects will affect the unfortunate recipient's upper torso.

Finally, fragmentation mines are designed to release fragments in a specific direction, usually effecting personnel at distances of up to 200 meters away from the mine emplacement area. These mines are particularly good for defensive protection of an area.

As stated earlier, AP mines tend to be significantly smaller than their AT cousins. This allows AP mines to be hidden more effectively, hence AP mines are typically more difficult to detect, generally speaking. Examples of AP mines are shown in Figure 3-12 and Figure 3-13, courtesy DoD and NPA.



Figure 3-12: Pressure activated AP mine, example #1



Figure 3-13: Pressure activated AP mine, example #2

Sizes range in the neighborhood of 6 to 12 cm in diameter and come in a variety of shapes.

3.4 Mine Emplacement

Mines can either be buried or laid on the surface of the ground. U.S. Army doctrine suggests that burial is the preferred method of emplacement, due to a higher degree of difficulty in detection, but any type of mine that is designed for burial can just as easily be laid on the surface. Quoting from the Army's Field Manual, "If time permits, mines should be buried to increase their effectiveness; but they can be laid on top of the ground." (Army, 2002) This creates a distinct challenge area in synthetic scene creation, as surface laid mines will have a different signature than buried mines. The details of mine signatures will be discussed in detail in a later section. To be effective, mines must be emplaced where they cannot be easily detected, but also in an area where a person or vehicle can apply enough pressure to detonate them. As stated before, AT mines are typically emplaced along roads, or in areas that vehicles will travel through. The term road is used loosely here, since tanks and other military vehicles are not bound by pavement or smooth paths. AP mines follow similar doctrine, but are targeted for areas that are heavily foot traveled.

The Army Field Manual lists very specific ways to emplace different types of AT and AP mines. Obviously, surface laid mines need only be set down and armed, but if they are to be buried, depth of burial is an important factor. Focusing on AT blast mines, 5 cm below the surface is considered the optimal burial depth. AP mines should be buried closer to the surface. Figure 3-14 shows a properly buried AT pressure mine.

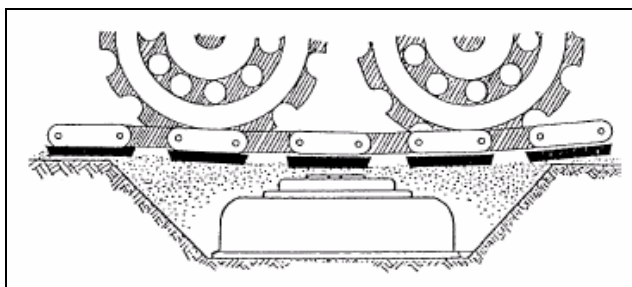


Figure 3-14: Employment of AT pressure mine (Army, 2002)

When burying landmines, care must be taken to ensure that the hole is large enough to allow sufficient weight to be exerted onto the surface of the mine. Therefore, holes are dug quite a bit larger than the mine itself. No specific doctrine for hole dimensions exist, therefore common sense is to be used. Concealment of the burial process is also paramount. This includes replacing sod or other altered surface features. A small mound should be left immediately after burial, as the soil will settle over time and a depression is unwanted. Figure 3-15 shows correct and incorrect mine burial techniques.

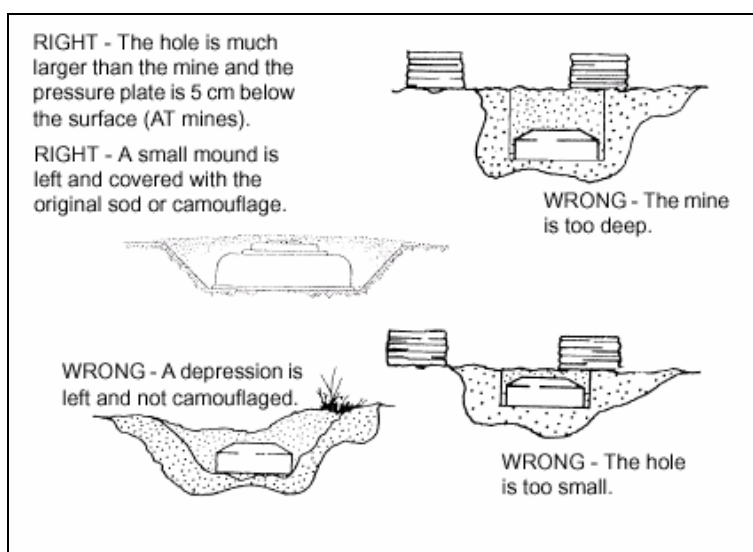


Figure 3-15: Proper mine burial techniques (Army, 2002)

3.4.1 Minefields

Emplacing more than one mine suggests a minefield, or specific areas in which groups of mines are employed or perceived to be employed (Army, 2002). Four types of minefields are noted in the Army Field Manual, protective, tactical, nuisance, and phony. Protective minefields are used in a defensive posture. They are classified as either hasty or

deliberate. Hasty implies that the minefield is temporary and laid in a short period of time. Surface-laid mines are frequent in this configuration. Additionally, low-metallic mines are not to be used, as this type of minefield is to be recovered once the emplacing force advances and low-metallic mines are difficult to remove. Mines are emplaced in rows with strict mine-laying doctrine upheld, an example of such a minefield is shown in Figure 3-16.

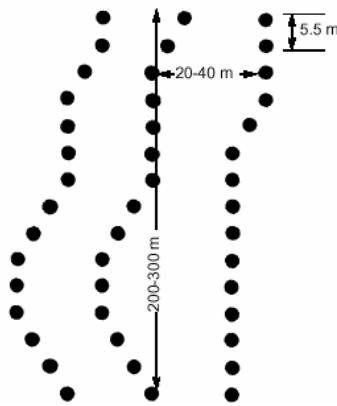


Figure 3-16: Standard rapidly emplaced row minefield (Army, 2002)

Deliberate minefields are used to protect strategic assets and are of a more permanent nature. A standard-pattern minefield is typically used, however a row minefield can also be emplaced. Intimate details of row and standard-pattern mining can be referenced in the Army Field Manual (2002).

Tactical minefields are laid in order to effectively control an advancing force and are of an offensive nature. These minefields can be emplaced in such a way that an advancing enemy is “funneled” into an area suitable for an ambush. These have structure similar to protective minefields, but located in a forward position. Nuisance minefields consist of irregularly placed mines or groups of mines. The nuisance minefield is typically the most difficult to detect in that any minefield structure is usually non-existent. Lastly, phony minefields are areas that are disguised to look like true minefields. Therefore, holes are dug, minefield markings are posted, and the area is treated like a true minefield. According to the Army, phony minefields are seldom employed without a true minefield nearby. The success of the phony minefield depends on the enemy’s state of mind, i.e. is the enemy mine-conscious?

3.5 Landmine Signatures

A landmine signature refers to the presence of a localized difference between a landmine and its surroundings caused by the landmine itself or the emplacement of the mine. This signature manifests itself differently in the visible or near infrared portion of the electromagnetic spectrum as compared to the MWIR or the LWIR regions. As this work focuses on the LWIR region for validation, the VIS/NIR region has not been emphasized. Signatures of both surface laid and buried mines are presented as potential detection features utilized by a detection algorithm. As the DIRSIG scene focuses on feeding algorithms, spectral signatures of landmines and their interaction with background objects are critically important. Validation of the scene is focused on this area.

3.5.1 Surface Landmine Signatures

The signature produced by a surface laid mine is directly due to the size, shape, composite material makeup, and thermal properties of the mine. These properties are inherently different than surrounding background objects such as soil, grass, etc. When viewed in the thermal region, these property differences will produce an apparent temperature contrast at the sensor. Algorithm developers use this contrast between the target mine and the background as a detection feature. To effectively model this contrast, detailed information must be known about the physical properties and spectral properties of the target mine as well as background objects. Additionally, potential false alarm targets will have similar, but different physical and spectral properties, producing target-like thermal contrasts. The ability of an algorithm to reduce false alarms depends on its ability to distinguish between subtle differences in thermal contrast, whether it be spectrally, spatially, or a combination of the two.

Two intertwined properties contribute to this apparent temperature difference, the emissivity and the temperature of the mine and the background. The emissivity of an object (for a particular wavelength and temperature) describes the amount of thermal energy radiated by the surface of the object. Emissivity can be described in the following manner mathematically (Schott, 1997).

$$\varepsilon(\lambda) = \frac{M_{\lambda}(T)}{M_{\lambda BB}(T)} \quad (3-9)$$

where, $M_{\lambda}(T)$ is the spectral exitance from an object at temperature T and $M_{\lambda BB}(T)$ is the spectral exitance from a blackbody at the same temperature. For opaque objects, Kirchhoff's law states that emissivity must be equal to absorptance for a surface in thermodynamic equilibrium (Schott, 1997). Applying Kirchhoff's law with conservation of energy, which states that all energy impinging on a surface must be either absorbed or reflected (for opaque objects) (Schott, 1997), we arrive at the following equation relating reflectance (r) to emissivity (ε) (Schott, 1997):

$$\varepsilon + r = 1. \quad (3-10)$$

The surface temperature of an object is a result of conduction, convection and radiation of heat energy through the object and with its ambient surround. Conduction refers to the transfer of heat energy by contact with another medium, such as the contact between a mine and the air around it or the mine and the ground it is laid on. Convection is the process by which heat flows across a gradient from a hotter point to a colder point, attempting to equalize within a given medium. Radiation is the process by which an object emits radiance. Self-emission of radiance at a wavelength of interest is determined through the Planck equation (3-5).

Radiance determined from Planck's equation is the total amount of thermal radiance that can be emitted from an object at a given wavelength and temperature. The object's emissivity can be thought of as a factor that allows only a certain percentage of that total thermal radiance to be emitted, for a given wavelength and temperature. This multiplicative combination of thermal radiance and emissivity comprises the D-type photons as shown in Figure 3-2. Total radiance reaching the sensor is determined through the "Big Equation", as described in equation (3-4). This, however, is not the same radiance that an imaging sensor would detect and used to develop a thermal contrast between targets and backgrounds. In order to develop the effective radiance sensed by the system, the spectral response function of the detector used to image the target must be incorporated (Schott, 1997). Specifically, a sensor observes radiance over a bandpass, or

range of wavelengths. The normalized spectral response function of the sensor describes how much of that total radiance reaching the sensor is observed by the detector at each wavelength over the sensor's bandpass. To determine the effective or observed radiance (L_{eff}), the following relationship applies (Schott, 1997).

$$L_{eff} = \int_0^{\infty} L_{\lambda} R(\lambda) d\lambda \quad (3-11)$$

where,

L_{λ} – output radiance from the “Big Equation”

$R(\lambda)$ – normalized spectral response function of the sensor

To be clear, equation (3-11) is presented in continuous form, as opposed to the discrete version, which is presented in equation (3-6). This radiance value described by L_{eff} is used to determine apparent temperature of objects and ultimately the thermal contrast between a target and the surrounding background.

It is difficult to generally characterize the thermal contrast between a mine and the background, as mines come in many different shapes, sizes and compositions. In addition, background variety can be seemingly endless to include wet soil, dry soil, short grass, tall grass, sand, etc. According to Nivelles and Lhomme (1997), soil type can play an important factor in the development of thermal contrast. They observed a contrast inversion when viewing a surface laid mine on rocky soil compared to the same mine laid on a compost background, all other factors held constant. In the first case, the mine was observed to have a negative contrast and in the latter a positive contrast at that particular time of day.

Another important consideration is the diurnal nature of the thermal contrast. This phenomenon is dependent on incident solar radiation as well as heat transfer due to conduction, convection, and radiation (Maksymenko, et al., 1995). Incident solar radiation contributes to heating the mine and background at different rates, depending on the thermal inertia of the mine and the background, as well as the emissivity of the surfaces of each. Maksymenko, et al. (1995) also point out that observed signatures are

far from constant, varying with atmospheric conditions. There are two noticeable crossover periods where the thermal contrast between the mine and the background is null. In general, these periods occur just after sunrise and just after sunset (Maksymonko, et al., 1995), but are heavily influenced by atmospheric conditions. Just after sunrise and assuming the mine in question heats or cools faster than the background, solar loading has warmed the mine from a point where the mine is cooler than the background to a point where the temperatures are equal. The mine should continue to heat faster than the background throughout the day, assuming a constant solar load. After sunset, solar loading has ceased. The mine and the background begin to transfer heat to the colder, nighttime sky. Again the mine will cool more rapidly than the background, reaching a point where temperatures equalize. From a detection standpoint, these times are not suitable, presenting a significant problem for 24-hour detection capabilities.

Numerous studies have been conducted over the past decade using different types of infrared imaging systems to observe the thermal contrast between specific surface mines and specific background environments over a diurnal cycle. Some papers include Janssen et al. (1996), Nivelles and Lhomme (1997), McGovern and Aponte, (2001), and Hong et al. (2002). Difficulty exists in directly comparing the results of the experiments due to such wide variety of environmental conditions, sensor types, backgrounds, and mine types. A conclusion that can be drawn from each of the studies is that the diurnal cycle is generally observable for AP and AT mines, either plastic or metallic, over a wide range of environmental conditions, given sufficient resolution and sensor sensitivity. Nivelles and Lhomme (1997) state, "Sensitivity of about 0.5K seems enough to see the mines with sufficient contrast and a spatial resolution between 0.5 and 1 inch seems enough to classify and detect any object as mine." At this resolution and imaging from an oblique angle, the authors were able to observe structural information about the observed AT mines, which they hypothesize can be used to help classify the object as mine by an automatic mine recognition program and distinguish between false targets.

A different approach to detect surface mines is presented by Cremer et al. (2002). They suggest that using linear polarization measurements of a surface laid landmine

scene may be able to detect the presence of landmines in heavy vegetation. Specifically their work dealt with polarization in the MWIR region. The basis for the approach lies with the known phenomena that radiation becomes partially linearly polarized when reflected or emitted from a smooth surface. Typically, man-made objects have smoother surfaces than natural occurring objects, so in theory, this approach will locate not only mines but other man-made objects within the scene of interest. The theory states that because the mines have a smoother surface than the surrounding natural objects, i.e. the soil, grass, etc., they should stand out when a linear polarized image of the scene is viewed. The potential for false alarms can be quite high with this approach if there are a number of man-made, smooth objects causing a linear polarization of the reflected and emitted radiation other than landmines. However, with some fore knowledge of the scene, false alarms could be reduced.

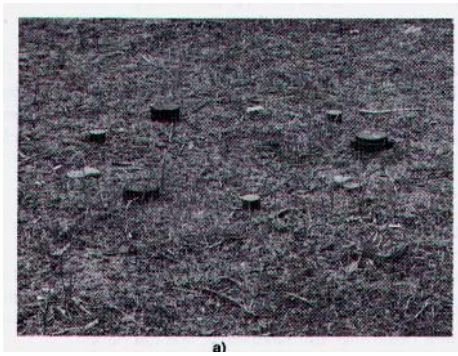


Figure 3-17: Polarization scene set-up

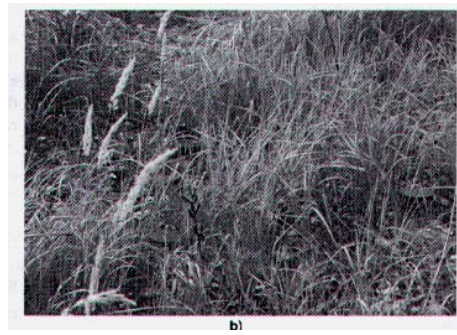


Figure 3-18: Imaged polarization scene after vegetation growth

In Figure 3-17 (Cremer et al., 2002), the scene of reference is shown with landmine-objects placed in a circular fashion. The same scene is shown overgrown with long grass in Figure 3-18. The results of viewing the overgrown scene with a linear polarization filter are shown in Figure 3-19 (Cremer et al., 2002).

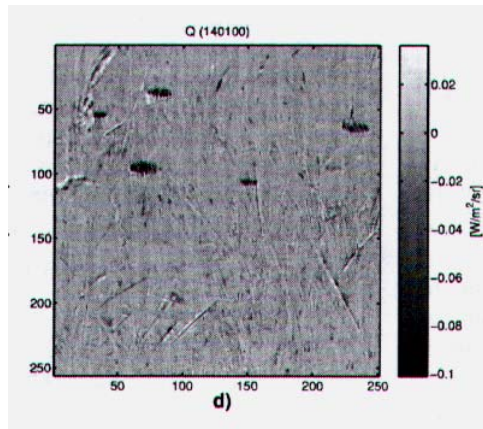


Figure 3-19: Linear polarization image

The use of polarization images to detect surface laid mines is a newer technology that is currently being exploited. It is presented as a reference to additional mine features that may be exploited by an algorithm, but has not been researched in this work.

3.5.2 Buried Landmine Signatures

A buried mine signature is different than that of a surface-laid mine in that the observable features are not of the mine directly, rather of what impact the mine has on the background. The discussion will be in two parts, the first discussing pure thermal effects observed by buried mines and the second discussing observed infrared spectral effects.

The observed thermal signature of buried landmines is an apparent temperature contrast between the surface temperature of the soil above the mine and the surface temperature of the soil surrounding the mine. There are two commonly observed thermal effects, namely the surface effect and the volume effect (Simard, 1996). The surface effect is associated with the process of disturbing the soil directly above the mine during the emplacement process. Disturbing the soil to emplace a mine causes a change in the density of the soil, such that it will have a lower thermal conductivity when compared to the surrounding undisturbed soil (DePersia et al., 1995). The lower thermal conductivity leads to a noticeable thermal contrast between disturbed and undisturbed soil. The surface effect is generally applicable to recently buried mines, as environmental conditions, such as rain or wind, will degrade apparent temperature contrast. According to Simard (1996), the effects can last up to three weeks under the right conditions. In

observing the surface effect, DePersia et al. (1995) suggest that a broadband infrared sensor may prove to be the best approach, mimicking the common observation technique for observing apparent temperature contrast of a surface-laid mine and its background.

The volume effect deals with the presence of the thermal mass of a buried mine in the soil. It is observed for as long as the mine is in the soil, but reduced depending on environmental and atmospheric conditions. The volume of soil directly above the buried mine will not heat up or cool down at the same rate as the surrounding soil, due to the influence of the mine's thermal mass. The effect is more pronounced depending on mine burial depth. The deeper the mine is buried, the smaller the apparent temperature contrast (Khanafer and Vafai, 2002). This observable temperature contrast at the surface of the soil follows a diurnal cycle similar to the diurnal cycle of a surface laid mine. It is noted through the work of McGovern and Aponte (2001) and Maksymenko et al. (1995) that the variation in apparent temperature contrast of a buried landmine over the diurnal cycle is less than that of a surface laid mine. Obviously the exact differences are influenced by mine type, burial depth, background type and other atmospheric parameters. Detailed thermodynamic analysis of a buried mine in soil is presented in Khanafer and Vafai (2002), where they use a finite element model to predict surface temperatures of soil directly over a buried mine through a diurnal cycle. In addition to determining the effect of burial depth, they also observed that variations in the surface, i.e. rough terrain versus smooth terrain produces significant variations in the observed thermal signature when compared to the results of a flat surface model. They suggest that failing to incorporate surface roughness when modeling this phenomenon may not adequately represent the occurrence of potential false alarm points in the data. The volume effect is also observed due to other objects, such as buried rocks, which may cause significant false alarms in a detection algorithm. The following example images from the work of Russell et al. (1997) show surface and volume effects.

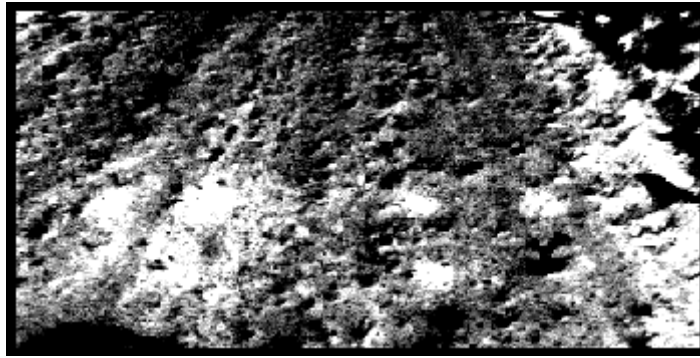


Figure 3-20: Daytime IR image of buried mines in soil

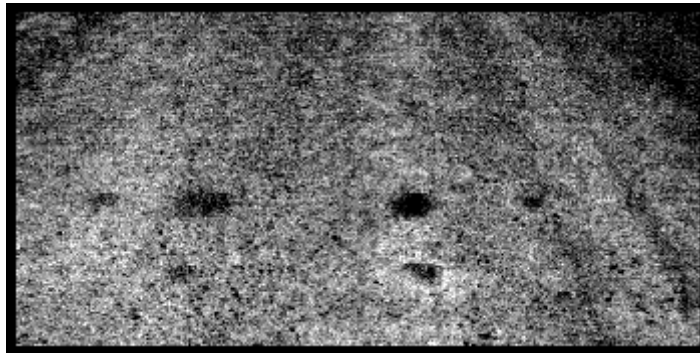


Figure 3-21: Nighttime IR image of buried mines in soil

In the images, the three mines on the left (two AP and 1 AT) have been buried only one week prior to imaging whereas the mines on the right have been buried for approximately one year. It is clear to see from these images the surface effect combined with the volume effect in the newly buried mines, especially in the daytime image. The thermal contrast in the long buried mines results directly from the volume effect alone.

The second part to buried landmine signatures relates to the spectral structure of the soil disturbed during the burial process and the undisturbed surrounding soil. Immediately after mine burial, the disturbed soil will exhibit a localized texture difference capable of being observed by a broadband IR sensor or even in the visible portion of the spectrum (DePersia et al, ISSSR 1995). However, detection based on this feature alone can result in false alarms due to naturally occurring texture differences. Additionally, changing environmental conditions such as heavy rainfall or blowing wind will effectively eliminate any observable texture difference due to mine burial. Therefore, localized texture differences are not the most reliable observable characteristic. According to DePersia et al. (1995), the main theory behind buried

landmine detection due to spectral properties of the soil depends on the difference in spectral structure between the surface layer of soil and the subsurface soil. Essentially during the burial process the subsurface soil is churned, resulting in some of the subsurface soil residing at the surface. A spectral difference can be observed between subsurface soil and surface soil based on weathering effects on the surface soil and organic composition differences. DePersia et al. (1995) suggest that this spectral difference is observable even after localized texture differences have been eliminated. The question remains, what kind of difference in the spectra of the two soil types is observed?

Over the MWIR and LWIR regions, soil will show spectral structure due to specific spectral features of the minerals contained within (Winter et al., 1996). Therefore, the simplest method of detection would be to observe a different spectral signature from the disturbed soil based solely on a change in mineral composition. If the disturbed and undisturbed soils are of equal mineral content, all hope is not lost. A spectral feature common to most soils is the Silicate Reststrahlen feature, which manifests itself in the 8.5 to 9.5 micron spectral window (Winter et al., 1996). This feature can be exploited to detect buried objects. An experiment by J. Salisbury as reported by DePersia et al. (ISSSR 1995) showed the effects of soil particle size on emissivity profiles in the LWIR. Before soil disturbance occurs, the subsurface soil layer is composed of large and small particles. Typically, large particles are covered by smaller particles. The surface soil layer, having been exposed to wind, rain, and other atmospheric effects, is only composed of large particles having been stripped of smaller particles. Salisbury measured emissivity spectra of large particle quartz and finer particle quartz and concluded that the Reststrahlen feature was much more pronounced for large particle quartz. DePersia et al. (ISSSR 1995) further argue that this observable difference in the reststrahlen bands should manifest itself in disturbed soil. They deduce that the disturbed soil will contain large particles covered by smaller particles, creating a higher observed emissivity in the reststrahlen bands as compared to the undisturbed surface soil, which contains only large particles. Experiments have been performed to verify this

hypothesis and are described in DePersia et al. (1995), DePersia et al. (ISSSR, 1995), and Winter et al. (1996). Each concludes that the hypothesis is valid. An example from Winter et al. (1996) shows the differences in emissivity of disturbed and undisturbed soil for a mine buried two weeks prior in a dirt road.

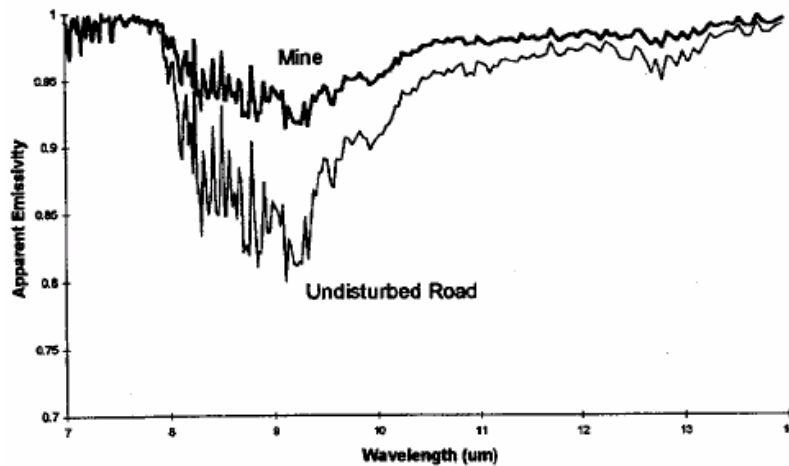


Figure 3-22: Spectral signatures of disturbed and undisturbed soil

As one may expect, the spectral difference has also been shown to decrease due to weathering. In an experiment by Winter et al. (1996), the buried minefield under investigation had received rainfall after burial. Additional to the rainfall, run-off rainwater from the surrounding area flooded the site, creating a true test of severe weathering. Spectra were recorded and the results are shown below.

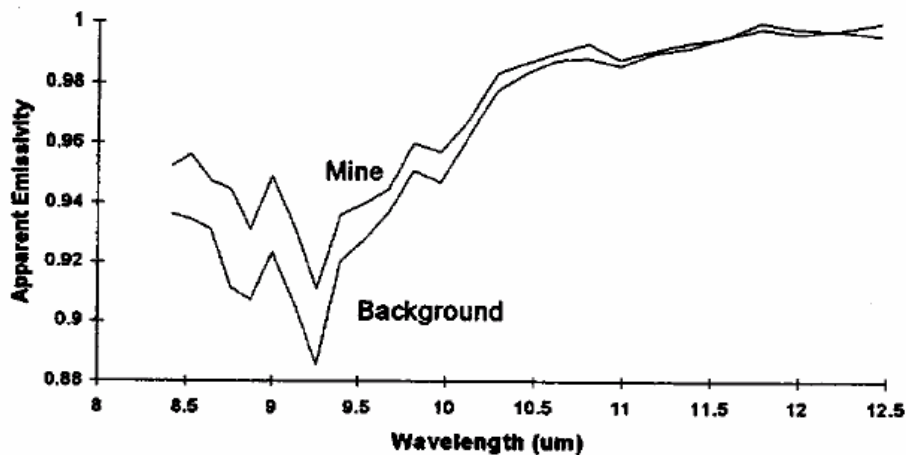


Figure 3-23: Spectral signatures for disturbed and undisturbed soil after flooding

Winter et al. (1996) state, “Note that the classic quartz feature is present in both the mine and in the background. The feature is stronger in the background location, although the difference is not as great as seen in the other detection examples.”

3.5.3 Signature Modeling

In order to develop a model that accurately represents a minefield, the incorporation of the thermal aspect and spectral aspect of a mine’s signature is of critical importance. From a thermal point of view, what distinguish a buried mined area from a non-mined area are the fundamental material properties of the soil directly above the mine from the soil around the mine. If the difference in these material properties are known, they can be modeled through the creation of two types of soil, each attributed with its own set of properties. The thermal model will develop the observed temperature difference between these soils in accordance with the rendered data’s time of day. From a spectral point of view, the distinguishing feature is simply a difference in emissivity between the two types of soil. This can be modeled in the same fashion as the thermal difference, through the use of two types of soil, each with its own associated emissivity curve, defining the observed spectral difference in the synthetic data. These concepts will be revisited in detail in Chapter 4.

3.6 Comparison Metrics

Once the synthetic landmine scene was built, comparison metrics were used to determine the accuracy of the data. These metrics compare the synthetic data to truth data that was collected over a mock minefield. This truth data will be discussed further in Chapter 4. The following section will address the theory behind each metric used to compare the synthetic scene to the truth scene. In particular, the metrics are Rank Order Correlation (ROC), Gray Level Co-occurrence Matrix (GLCM) analysis with emphasis on the contrast metric derived from the matrix, Spectral Co-occurrence Matrix (SCM) analysis with emphasis on the contrast metric derived from the matrix, and the R(x) anomaly detection algorithm. Specific application of each metric to the data used in this work is discussed further in subsequent sections.

3.6.1 Rank Order Correlation

The synthetic scene has been developed with the goal of training and assessing ATR algorithms, focused on landmine detection. Therefore, the thermal contrast between mine-like targets and the background is a detection feature that must be compared for accuracy. For the type of analysis needed in this work, exact temperature differences between synthetic and truth pixels will be of little benefit, as a mine detection algorithm focusing on contrast will be insensitive to a scene-wide gain or bias change in temperature (Mason et al., 1994). With this in mind, a metric of Root Mean Square (RMS) error will be insufficient. Instead, a measure of relative contrast between objects within the scene would be more appropriate. An example in Mason et al. (1994) explains this quite well. They state, “If the model introduced a constant bias error of 5°C when expressed as apparent temperature, the RMS error would be 5°C. However, a contrast based algorithm would consider either image equally acceptable since the relative contrast rank of each object remains essentially the same.” The objective is to quantify the ability of DIRSIG to model the relative brightness of objects within the scene, preserving the overall brightness ranking of feature objects. Particularly, as the contrast will change over the diurnal cycle, measurements need to be of multiple times of day so that proper characterization of contrast over the diurnal cycle can take place.

Rank Order Correlation is used to measure brightness rankings of feature objects between truth and synthetic imagery at different times of day. As defined in Mason et al. (1994), Spearman rank order correlation is as follows:

$$\rho_{\text{TOD}} = 1 - \frac{6 \sum (R_i - R'_i)^2}{n^3 - n} \quad (3-12)$$

where,

ρ_{TOD} – correlation coefficient for each image pair at a specific time of day

n – number of samples

R_i – rank in the truth image for the i^{th} object

R'_i – rank in the DIRSIG image for the i^{th} object

The application of ROC is straightforward. Objects of interest covering the scene's range of brightness are ranked in each scene, from brightest to darkest. Each object's ranking and the number of comparison point pairs are inserted in the ROC equation, generating a correlation coefficient for that scene combination at a particular time of day. The correlation coefficient statistic runs from -1.0 to 1.0 , where 1.0 is perfect correlation. Plotting rank order in the truth image versus rank order in the DIRSIG image will give an initial evaluation of the data (Kraska, 1996). A perfect match between scenes would produce a linear graph. For diurnal studies, plotting correlation coefficients versus time of day can also point out modeling discrepancies.

3.6.2 Gray Level Co-occurrence Matrix

Evaluating the synthetic scene's ability to statistically represent the truth image is the reason for employing the Gray Level Co-occurrence Matrix (GLCM) approach. As most ATR algorithms are statistically based, this evaluation will be invaluable in assessing the statistical fit between the two scenes. The GLCM captures the spatial distribution of gray level values within an image or a subset region over a single bandpass. Specifically, the GLCM is comprised of entries that represent the relative frequency of gray value co-occurrences between a pixel and its nearest neighboring pixels. Once this matrix is developed, statistical measures of the textural content of the image can be derived utilizing the GLCM entries. Haralick et al. (1973) discusses 14 textural feature metrics obtained through the GLCM, however this discussion will be limited to six main features.

For description purposes, the neighbors of a pixel of interest are described by Haralick et al. (1973) using distance (d) and four angular (α) relationships. The left and right adjacent pixels are $d=1$, $\alpha=0^\circ$ neighbors, the upper and lower adjacent pixels are $d=1$, $\alpha=90^\circ$ neighbors, the upper right and lower left adjacent pixels are $d=1$, $\alpha=45^\circ$ neighbors, and the upper left and lower right adjacent pixels are $d=1$, $\alpha=135^\circ$ neighbors. Four versions of the GLCM can be calculated, one for each angular presentation while holding pixel distance constant at 1.0 . The GLCM is a symmetric $m \times m$ matrix where m

represents the total number of gray level values within the image. The GLCM is represented by Haralick et al. (1973) as $P(i,j,d,\alpha)$ where i and j represent individual gray level values such that $m \geq i \geq 0$ and $m \geq j \geq 0$. An explanation through simple example is the best way of explaining how the $P(i,j)$ values are calculated. From Haralick et al. (1973), we will consider developing $P(i,j,1,\alpha)$ for a small 4×4 pixel image.

0	0	1	1
0	0	1	1
0	2	2	2
2	2	3	3

Figure 3-24: GLCM example image

Gray Levels	Gray Levels			
	0	1	2	3
0	#(0,0)	#(0,1)	#(0,2)	#(0,3)
1	#(1,0)	#(1,1)	#(1,2)	#(1,3)
2	#(2,0)	#(2,1)	#(2,2)	#(2,3)
3	#(3,0)	#(3,2)	#(2,3)	#(3,3)

Figure 3-25: General GLCM form for image with four gray levels

The general form for the GLCM of any image with $m=4$ is shown in Figure 3-25 where (#) refers to the number of co-occurrences between gray levels. To populate the matrix with entries in the $\alpha=0^\circ$ case, the number of times the horizontal combination of gray level values shown in the template matrix is counted for each combination possible. This is done left to right then right to left. Examining the entry for (0,0), we first look at each pixel and compare it with it's neighboring pixel to the right and add the number of zero followed by zero occurrences. We can see two times where there is a (0,0) left to right combination. Now, examining each pixel for a (0,0) right to left combination, we see two more. Therefore the (0,0) entry for the $\alpha=0^\circ$ case is 4. The resulting matrices are shown in Figure 3-26 for all four angular presentations.

$$P(i,j,1,0) = \begin{bmatrix} 4 & 2 & 1 & 0 \\ 2 & 4 & 0 & 0 \\ 1 & 0 & 6 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix} \quad P(i,j,1,45) = \begin{bmatrix} 4 & 1 & 0 & 0 \\ 1 & 2 & 2 & 0 \\ 0 & 2 & 4 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

$$P(i, j, 1, 90) = \begin{bmatrix} 6 & 0 & 2 & 0 \\ 0 & 4 & 2 & 0 \\ 2 & 2 & 2 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix} \quad P(i, j, 1, 135) = \begin{bmatrix} 2 & 1 & 3 & 0 \\ 1 & 2 & 1 & 0 \\ 3 & 1 & 0 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix}$$

Figure 3-26: GLCM example matrices

The GLCM matrices serve as the basis for obtaining the desired statistical information, namely the statistical representation of the gray level variation within the image. As alluded to earlier, 14 statistical metrics can be derived from the matrix as defined by Haralick et al. (1973), six of which will be presented here. Inherent to all metrics derived from the GLCM, $P(i, j, d, \alpha)$ must be normalized by the number of gray level values in the image, such that $p(i, j, d, \alpha)$ is the normalized GLCM given by,

$$p(i, j, d, \alpha) = \frac{P(i, j, d, \alpha)}{m} \quad (3-13)$$

Using the normalized GLCM, metrics of Angular Second Moment (Energy), Contrast, Correlation, Variance, Inverse Difference Moment, and Entropy can be defined. (Haralick et al., 1973)

1) Angular Second Moment (Energy)

$$f_1 = \sum_{i=1}^m \sum_{j=1}^m [p(i, j)]^2 \quad (3-14)$$

2) Contrast

$$f_2 = \sum_{n=0}^{m-1} n^2 \left[\sum_{i=1}^m \sum_{j=1}^m p(i, j) \right], \text{ where } |i-j| = n \quad (3-15)$$

3) Correlation

$$f_3 = \frac{\left[\sum_{i=1}^m \sum_{j=1}^m (i \cdot j) \cdot p(i, j) - \mu_x \mu_y \right]}{\sigma_x \sigma_y} \quad (3-16)$$

where, μ_x , μ_y , σ_x , and σ_y are the means and standard deviations of the rows and columns of the matrix P.

4) Variance

$$f_4 = \sum_{i=1}^m \sum_{j=1}^m (i - \mu)^2 \cdot p(i, j) \quad (3-17)$$

5) Inverse Difference Moment

$$f_5 = \sum_{i=1}^m \sum_{j=1}^m \frac{p(i, j)}{1 + (i - j)^2} \quad (3-18)$$

6) Entropy

$$f_9 = - \sum_{i=1}^m \sum_{j=1}^m p(i, j) \cdot \log[p(i, j)] \quad (3-19)$$

This work uses the contrast metric, as the metric measures the amount of local variations in the image. This metric characterizes the contrast over a specified subset of the imaged scene, known to contain landmine thermal signatures. If the contrast values match reasonably well between DIRSIG and truth imagery then we can conclude that statistically, these scenes match well. For accurate comparison purposes, the DIRSIG scene and the truth data should be spatially equivalent, ensuring that spatial variations in gray level are characterized over the same physical area. However, this may not be practical in this evaluation as perfect geometric accuracy is not required. Therefore, general similarity between the contrast values calculated for DIRSIG and truth imagery is all that can be expected.

3.6.3 Spectral Co-occurrence Matrix

Another approach to defining the statistical fit between the two scenes relies on the Spectral Co-occurrence Matrix (SCM) approach. Similar to the GLCM approach, a co-occurrence matrix is created corresponding to relative frequencies of gray level values of neighboring pixels. From this matrix, the same statistical metrics as defined in the GLCM section can be determined. Apart from the GLCM, the co-occurrence matrix in SCM calculations is determined across two different spectral bands, creating a statistical measure of spectral quality. This approach is detailed in Scanlan (2003) as a quantifiable measure of texture quality between spectral bands. As with the GLCM, an explanation through simple example is the best way of explaining the calculation of the SCM. The

following example is from the ENVI User's Guide (1999). We will consider developing the SCM for a 3×3 pixel processing window over an image containing two spectral bands of data. The “base” processing window is shown in Figure 3-27 and is taken from the first spectral band.

4	3	5
3	5	6
6	4	3

Figure 3-27: SCM base window from band #1

3	5	6
5	6	3
3	4	6

Figure 3-28: SCM shift window from band #2

The base window is theoretically overlaid on a shift window determined by angular presentation and distance, exactly the same as in GLCM calculations. In this example the distance is 1.0 and the angle is 0°. This corresponds to an X direction shift of one pixel and a Y direction shift of zero pixels. If the shift window values are obtained from the same spectral band as the base window, a GLCM is calculated, however, if the shift window values are obtained from a different spectral band, an SCM is calculated. In this example the shift window values are obtained from spectral band #2 (as shown in Figure 3-28), therefore an SCM is calculated. Mentally overlaying the base window onto the shift window will show the co-occurring entries for the matrix, as shown in Figure 3-29.

		<u>Gray Levels</u>			
		3	4	5	6
<u>Gray Levels</u>	3	0	0	2	1
	4	2	0	0	0
	5	0	0	0	2
	6	1	1	0	0

Figure 3-29: Resulting spectral co-occurrence matrix

From this point, contrast or any of the other metrics derived from this matrix are calculated exactly as discussed in section 3.6.2. In summary, the SCM is a variant of the GLCM that is developed over two different spectral bands of an image rather than within a single spectral band. The use of this evaluation will provide insight into the spectral accuracy of the synthetic image.

3.6.4 R(x) Anomaly Detection Algorithm

The final measurement of comparison between the DIRSIG and truth data is the evaluation of an anomaly detection algorithm. The focus here is not to develop or evaluate the capabilities of the algorithm to find mines, rather to ensure the algorithm performs similarly on both sets of data. An anomaly detector simply attempts to identify areas that may be out of the ordinary from the local surround. No specific knowledge about the target is needed. One of the most widely accepted and utilized algorithms for this type of work is the R(x) algorithm, developed by Reed and Yu (1990). R(x) is unique in that it encompasses not only spectral information but spatial information as well. The following derivation and explanation of the R(x) algorithm follows closely with the treatment in Schott (2003).

Let a pixel vector \mathbf{x}_n be composed of digital count values across j spectral bands such that,

$$\mathbf{x}_n = \begin{bmatrix} DC_1 \\ DC_2 \\ \vdots \\ DC_j \end{bmatrix} \quad (3-20)$$

Then, a sub-image comprised of n pixel vectors is selected to form \mathbf{X} , a $j \times n$ matrix. For example, if the region of interest was over four pixels spatially distributed as follows,

\mathbf{x}_1	\mathbf{x}_2
\mathbf{x}_3	\mathbf{x}_4

the resultant matrix \mathbf{X} will look like,

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4] \quad (3-21)$$

where each \mathbf{x}_n is a j -dimensional vector.

Similarly, a target's spatial shape over the n pixels in the sub-image is denoted by \mathbf{s}^T , a $1 \times n$ vector,

$$\mathbf{s}^T = [s_1, s_2, \dots, s_n] \quad (3-22)$$

where the subscript n represents spatially the same n pixels as represented in \mathbf{X} . Each entry in \mathbf{s}^T is proportional to the fractional content of the target in each pixel where,

$$\mathbf{s}^T \mathbf{s} = 1 \quad (3-23)$$

Additionally, a known target vector of dimension $1 \times j$ is denoted by \mathbf{b}^T where,

$$\mathbf{b}^T = [b_1, b_2, \dots, b_j] \quad (3-24)$$

In order to apply \mathbf{b}^T and \mathbf{X} in the algorithm, each pixel vector must be de-meanned by subtracting the mean value of all pixels in the sub-image from each \mathbf{x}_n and \mathbf{b}^T vector.

Following demeaning, the local covariance \mathbf{M} is computed. \mathbf{M} is a $j \times j$ matrix calculated over the sub-image as follows,

$$\mathbf{M} = \frac{1}{n} \mathbf{X} \mathbf{X}^T \quad (3-25)$$

where n represents the total number of pixel vectors in the sub-image. The constant false alarm rate detector (CFAR) version of the $R(x)$ algorithm is now presented (Schott 2003).

$$r(\mathbf{X}) = \frac{(\mathbf{b}^T \mathbf{M}^{-1} \mathbf{X} \mathbf{s})^2}{\left[1 - \frac{1}{n} (\mathbf{X} \mathbf{s})^T \mathbf{M}^{-1} (\mathbf{X} \mathbf{s})\right] [\mathbf{b}^T \mathbf{M}^{-1} \mathbf{b}]} \begin{matrix} H_0 \\ \geq r_0 \\ H_1 \end{matrix} \quad (3-26)$$

Where r_0 is the user determined threshold value such that if the $r(\mathbf{X})$ value is greater than or equal to r_0 then we accept the H_0 hypothesis, that there is target present. Otherwise we accept the H_1 hypothesis, that there is no target present.

Rather than utilizing the target detector, an adaptation of the CFAR detector designed for anomaly detection is of interest in this work. Therefore, to make the change, a scene-derived estimate of the target's spectral signature must be substituted for \mathbf{b} and \mathbf{b}^T . In this case the best estimate for \mathbf{b} is given as $\mathbf{X} \mathbf{s}$ (Schott, 2003). Using $\mathbf{X} \mathbf{s}$ for \mathbf{b} and substituting in for \mathbf{M} based on equation (3-25), the following CFAR anomaly detector version of the algorithm is obtained (Schott, 2003).

$$r(\mathbf{X}) = \frac{n(\mathbf{Xs})^T (\mathbf{XX}^T)^{-1} (\mathbf{Xs})}{1 - (\mathbf{Xs})^T (\mathbf{XX}^T)^{-1} (\mathbf{Xs})} \begin{matrix} H_0 \\ \geq \\ H_1 \end{matrix} r_0 \quad (3-27)$$

The $r(\mathbf{X})$ value is compared to a threshold value in the same manner as in the target detector version of the algorithm. These pixels that surpass that threshold value can be displayed in an image, showing where anomalies exist within the data.

As will be discussed in detail in Chapter 4, the synthetic scene created for this work was not built to match the truth scene pixel-for-pixel. Because the data are not perfectly registered and matched one-for-one, the expectation of the $R(x)$ algorithm's results cannot be the same as if the data had been perfectly matched. Due to the inclusion of a spatial representation of the target in the algorithm, some discrepancy between resulting $r(X)$ values in the truth and synthetic data is expected. In addition, differing image sizes between the truth and synthetic data poses an additional problem for a purely quantitative evaluation of the algorithmic results. Given the expected differences in the data, the algorithmic results can only be used to suggest similar scene structure from a spectral and spatial point of view. Implementation details of the $R(x)$ algorithm will be discussed in Chapter 5.

4.0 Approach

This section begins with a description of the Microscene experimental data collection performed on RIT's campus, August 2003. Next will be a description of the US Army Arid test site experimental collection, the site upon which the synthetic scene is based, as well as a description of the ground truth and other available data collected during the experiment. Following this will be a thorough description of the approach taken to develop the synthetic scene.

4.1 Microscene Experimental Data Collection

Microscene is an ongoing DIRS group project, which consists of creating a high-resolution (approximately 3-inch) scene based off a 150 by 150 meter geographic area on RIT's campus, depicted in Figure 4-1.



Figure 4-1: Microscene area

The purpose of the project is to use this scene as a “laboratory” for the introduction and evaluation of high-resolution scene modeling techniques in DIRSIG. As can be seen from Figure 4-1, Microscene encompasses numerous modeling challenges associated with high-resolution scene construction. From the geometric complexity to the wide

variety of vegetation and man-made objects within Microscene, this area allows all these challenges to be addressed. To facilitate the construction of a test bed of this magnitude, an extensive ground-truthing experiment was planned and executed. At the onset of this experimental description, the reader is also referred to Barcomb (2004) for additional information. The experimental planning and execution was a shared effort between the author and numerous DIRS Staff members. Since the focus of this work is to detail the data collected for investigation of buried and surface mine phenomenology, the concealed and camouflaged target aspect of the experiment is not emphasized. The hope is to create, between the two theses, a comprehensive overview of the experiment that can be used as a guide for those that may use the data in the future.

The overarching goal of the experiment was to capture representative signatures of surface and buried mine surrogates, concealed and camouflaged targets, and background objects, covering the spectrum from the visible to the long-wave infrared. To do this, RIT's Modular Imaging Spectrometer Instrument (MISI), a hyperspectral line-scanning sensor, and the Wildfire Airborne Sensor Program's (WASP) visible and infrared cameras were used to image the Microscene area. Placed within the area were surface and buried mine surrogates as well as concealed and camouflaged targets, all set-up according to known doctrine. An overview of the experimental set-up is shown in Figure 4-2.



Figure 4-2: Microscene experimental set-up

MISI has five LWIR multispectral channels, two 8-14 μm channels, one 8-10 μm , one 10-12 μm , and one 12-14 μm , along with 72 VIS/NIR channels (MISI webpage, 2003). The WASP system uses four separate broadband cameras, one 16-megapixel camera covering the 0.4-0.9 μm region, a 640×512 pixel SWIR camera covering 0.9-1.7 μm , a 640×512 pixel MWIR camera covering 3.0-5.0 μm , and a 640×512 pixel LWIR camera covering 8.0-9.2 μm (McKeown, 2003). These imaging systems were hoisted above the experimental area through the use of a scissors-lift to a height of approximately 50 feet above ground level (Figure 4-3). This placed the cameras approximately 125 feet from the center of the mine area. The WASP thermal cameras have an Instantaneous Field of View (IFOV) of 1-milliradian producing a resolution of approximately 1.5 inches at that distance. The WASP visible camera produced resolution of approximately 0.25 inches. MISI's IFOV is reported as 2 or 3-milliradians, producing a resolution of approximately 3 inches (at 3 milliradians).



Figure 4-3: Raised scissor cart

In the mine area, surface and buried mine targets were placed or buried within an area of soil mostly void of vegetation growth. This gave a consistent sandy background for which to image the thermal contrast exhibited by these targets over a diurnal cycle. This area is shown in Figure 4-4, to the left of the shed. Eight separate mine surrogates were used during the experiment, four buried and four surface-laid. Due to the unavailability of true surface mines, objects of similar dimensions were used instead. Two round weight disks of different sizes from a home gym, a rectangular piece of concrete and a cylindrical tin, spray-painted an olive-green color, were imaged and measured as representative surface mine targets throughout the experiment. The first weight disk measures approximately 27.5 cm in diameter, 6.5 cm thick and weighs 8 kg. The disk is covered with a gray plastic coating (covering thickness of ~2-3 mm) and contains a concrete core.



Figure 4-4: Mine area in Microscene

The second weight disk is similar to the first, containing a concrete core but covered with black plastic (covering thickness of ~2-3 mm) with a diameter of 24 cm, thickness of 7.0 cm and weight of 6.8 kg. The concrete block has outer dimensions of 33 x 27.5 x 4.5 cm weighing 10.9 kg. The painted tin cylinder has a diameter of 19 cm, height of 6.5 cm and weighs 2.25 kg. The cylinder was filled prior to the experiment with the same soil that the surrogates were laid upon for the inclusion of some representative thermal mass. Each of these surface surrogates has been geometrically recreated in .obj and .3dm formats and are available for use in DIRSIG.

The buried surrogate mines used in the experiment are SIMS, Simulant Mines, developed by the US Army for buried mine data collection. The SIMS test set is shown in Figure 4-5. According to the Army's Countermine division, SIMS are produced to accurately represent the thermal properties of live landmines (SIMS brochure, 2003). The SIMS test set has been graciously donated by Mr. Richard Ess, U.S. Army program manager for SIMS development, for use during this experiment. Mr. Ess has also provided technical documentation describing the material makeup and physical properties of the individual simulants. The four largest sizes (in diameter), 30 cm, 25 cm, 20 cm, and 12 cm SIMS were buried in the scene at a burial depth of approximately 5 cm. This follows known doctrine as outlined in Section 3.0.



Figure 4-5: SIMS test set, courtesy SIMS homepage

The mine surrogates were distributed in the scene in a manner that minimized insolation differences due to shadowing over the course of the day. Additionally, the spacing between mines was selected to minimize thermal influence of the targets on each other and to minimize the area needed to contain the surrogates. The layout is depicted in Figure 4-6.



Figure 4-6: Mine area layout

The buried mines are located beneath the multi-colored flags, closer to the blue control tarp. Each mine surrogate is placed 5 feet from the neighboring target and the two rows are 5 feet apart. A more detailed layout is shown in Figure 4-7.

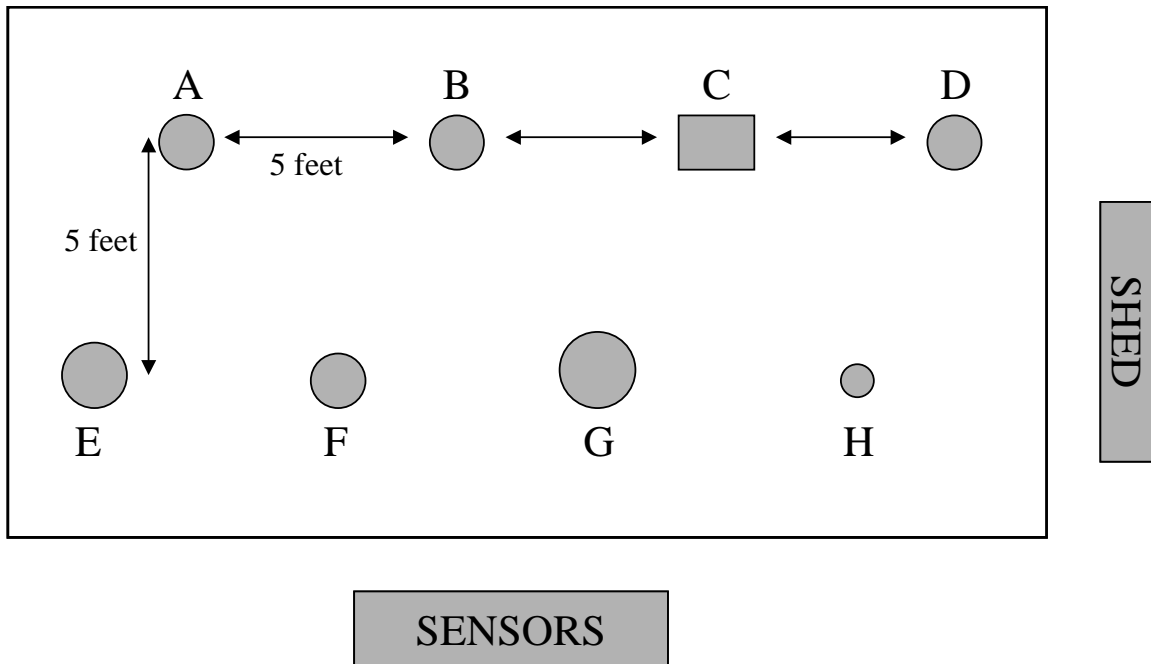


Figure 4-7: Detailed mine area layout

In Figure 4-7, A is the black disk, B is the tin cylinder, C is the concrete block, D is the gray disk, E is the 25 cm SIM, F is the 20 cm SIM, G is the 30 cm SIM, and H is the 12 cm SIM. The position of the shed and sensors is shown for relative positioning in the scene.

In addition to collecting data from the MISI and WASP systems, extensive ground truth information was obtained. Specifically, current weather information was collected through the use of a portable weather station capable of recording wind speed, temperature, and relative humidity. Additional radiosonde data was obtained from the closest available ground station in Buffalo, NY. While this is not ideal, it is the best available. Total direct and downwelled irradiance was collected during the experiment as well as temperatures of control and target objects within the scene. Each surface target was outfitted with a contact thermocouple while the buried ground areas were measured with staring IR thermocouples for truth temperature recording. Additionally, two staring IR thermocouples were used to obtain background surface temperature measurements of the sandy soil. These, along with other temperature readings are listed in the TEMP

section of the data set. For radiometric calibration purposes, two blackbodies, one CI Systems IR Radiation source Controller and one Techne TU-16A, both set at 40°C were placed in-scene. Crucial to the development of this experimental set-up in DIRSIG is the need for emissivity measurements of background objects and feature objects. Emissivity measurements of surface targets were made prior to emplacement in the scene using the SOC 400T FTIR reflectometer and the ASD spectroradiometer. The combination of these two measurements allows full spectral coverage from 0.4 to 25 microns.

As this collection was designed to obtain data for numerous audiences and uses, control tarps, control panels, a running generator, US military Humvees, and 2 metal panels with and without IR resistant coating were imaged along with the mined area. The next subsection deals with the timeline of events for this massive experimental activity.

4.1.1 Timeline of Events

This section is designed to give the reader an account of the steps that were taken to logistically coordinate the experiment, should someone need to undertake an effort like this in the future.

A number of weeks prior to the experiment (2+), a small committee of individuals assembled to lay the groundwork for the planning and execution of the collect. Monday August 25, 2003 was selected as the day to begin the experiment. This date coincided with the availability of the two imaging systems and was before the color change of the fall foliage. Plus, students were not yet on campus for the start of the fall quarter, allowing the experimental team more access to the area than may normally be allowed. Since data was to be collected over a 24-hour period, at the top of every hour, good weather was needed for at least this length of time. To ensure that 24 consecutive hours of good weather was obtained, a week timeframe was selected for the collect. Long-lead items were also identified and set into motion. Specifically, the need to borrow Humvees and surrogate landmines, rental of the scissor-lift, and the adaptation of MISI to a suitable configuration were addressed. MISI was to be placed on its side and outfitted with a

motorized turntable to simulate aircraft movement in the along-track direction. The scan mirror scanned the scene in the vertical direction and the turntable swept out lines in the horizontal direction. This MISI configuration had not been utilized previously and required additional time to get ready. Additional long-lead items included equipment to be purchased. In this case, we needed to buy camouflage netting, thermocouples, a data logger with multi-channel capability, and the associated cables and wires for temperature monitoring. Weekly meetings were scheduled for follow up on the coordination.

One week prior to the experiment we coordinated with RIT's Facilities Management. From Facilities Management we were able to obtain a 21-kilowatt generator filled with gasoline, which was capable of powering all required computer equipment for both imaging systems. Additionally, Facilities Management provided a canopy tent used to cover all the computer equipment, tables, chairs, and related equipment. All-weather tarps were obtained and kept on hand for briefly covering the imaging systems in case of adverse weather. In addition to Facilities Management, we coordinated with the Center for Industrial Management Systems (CIMS) and RIT's building 7 personnel for use of their high bays. During periods of adverse weather, the imaging systems were protected by driving the scissor cart into the high bay at building 7. One week out, a shift schedule was created, ensuring adequate personnel coverage during the entire experiment. Excess debris left in the scene from the archery team was removed and notification was made to Campus Safety informing them of the experiment. Final preparation was made to use the imaging systems remotely on the ground rather than staffing a person on the lift with the systems during the experimental period.

The day prior, the buried mine surrogates were emplaced. This ensured thermodynamic equalization occurred prior to imaging the area. The morning of the experiment, all target and data-logging equipment were installed. Miscellaneous items procured the day of included, liquid nitrogen for MISI's dewers, clothing, entertainment, and food items for personnel.

4.1.2 Collected Data

Data was collected with both systems at the top of each hour. Weather and temperature information was collected and recorded in 5-minute increments for the duration of the experiment. First images were collected at 1900 hrs on Aug 25, 2003. Hourly images were collected through the night and halted after the 0600 images due to rain in the area. At this time the emergency “adverse weather” contingency plan was put into action, requiring the movement of the systems from the experimental area to the high bay at building 7. The weather cleared and imaging was resumed on Aug 27, 2003 at 1000 hrs. Good, clear weather lasted for the duration of the experiment. From 1000 hrs on Wednesday, through 1500 hrs on Thursday, images were collected at the top of every hour with both imaging systems. This resulted in 43 collections over the course of the experiment. A few example images are shown in Figure 4-8 through Figure 4-10.



Figure 4-8: WASP Visible at 1400hrs

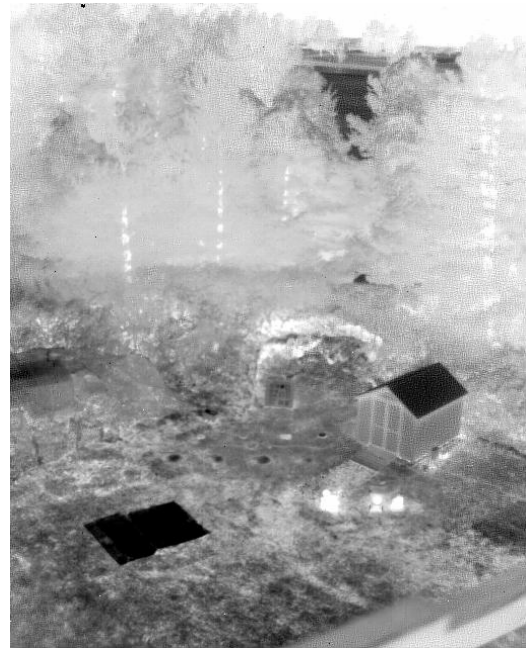


Figure 4-9: WASP LWIR at 0200hrs



Figure 4-10: MISI VIS Image, 1000hrs

Extensive analysis of the data obtained from this experiment has not been accomplished within this body of work as the modeled scene stems from a different data set. From initial observation, the imagery taken from the WASP system seems excellent. MISI imagery seems to be somewhat suspect at certain times of acquisition. There appears to be glint in the imagery, stemming from the accidental inclusion of stray light. The reader is referred to Barcomb (2004) for further detail on this subject. All data collected from this experiment is located on the DIRS static drive in the /cis/static/dirs/MicroScene directory. Included with the collected data on the static drive is a PowerPoint briefing detailing the intricacies of the experimental planning process. At the time of this writing, the spectra collected of individual targets have not been uploaded to the static directory. A discrepancy has been noted with the collection methods used to obtain curves, however it is hoped that not all of the spectral ground truth is unusable. The DIRS group measurements team is reviewing the curves and should upload good data to the static drive after the review is completed.

The author hopes that this collection can be used as a model for future experiments. We had remarkably few complications and would have had less given more experience in planning a scientific event of this nature.

4.2 US Army Arid Test Site Scene

As discussed previously, the scene created in DIRSIG has been based on the US Army Arid test site data collect performed at the beginning of April 2003 in Arizona. The area selected for modeling is one of three sites imaged during the collect. The particular site selected was chosen due to the better availability of ground truth. This region contains a surface and buried minefield, laid in accordance with known doctrine (see Figure 4-11).



Figure 4-11: Collection area

In the scene are four main target types, plastic anti-tank mines, metallic anti-tank mines, Electro-optical (EO) calibration targets (red paper-like squares), and Top Hat fiducial targets. The plastic and metallic mines are found along the surface as well as buried flush to the ground within the scene. The square mines on the right in Figure 4-11 are plastic mines and the round objects on the left are metallic mines. An EO target is the first, red object on the left. The Top Hat targets are shown in Figure 4-12 as light colored, tall, cylindrical objects.



Figure 4-12: Top Hat fiducial target

According to members of the Army MURI team, the Arid test site has been imaged by at least two systems, the U.S. Navy's LASH system and the University of Hawaii's AHI system. The Littoral Airborne Sensor Hyperspectral (LASH) system is a hyperspectral pushbroom imager that covers the VIS/NIR portions of the spectrum (Stein et al., 2001). Since this work deals directly with the LWIR, the data collected from LASH of the site was not heavily exploited.

AHI, Airborne Hyperspectral Instrument, maintained by the University of Hawaii, is also a hyperspectral pushbroom instrument. AHI covers the LWIR portion of the spectrum from $7.5\mu\text{m}$ to $11.7\mu\text{m}$ using a 256×256 HgCdTe focal plane array. The spectral resolution of the system is 100 nm (AHI Homepage, 2003). Data from the site, taken from AHI is the backbone for site creation in DIRSIG and is utilized for the final quantitative and qualitative comparisons. AHI configuration will be discussed further in later sections.

4.2.1 Available Ground Truth

At the outset of this project very little was known about the ground truth information that was collected during this experiment. Since the main feature that an algorithm would be triggering off for buried mine detection is the spectral Reststrahlen

feature, the author hoped for some extensive truth emissivity curves taken on and off some buried mine areas. Unfortunately, this was not available. In fact, ground truth suitable for modeling was limited in general. Once the author was granted access to the FTP site, which housed the data collected during the experiment, at least two solid weeks were spent sifting through the information trying to determine what data were useful. A catalog of 24 ground level photos, similar to Figure 4-11 were obtained. Along with the images is a detailed listing of where each image was taken and the direction in which the photographer was facing. This is shown in Figure 4-13.

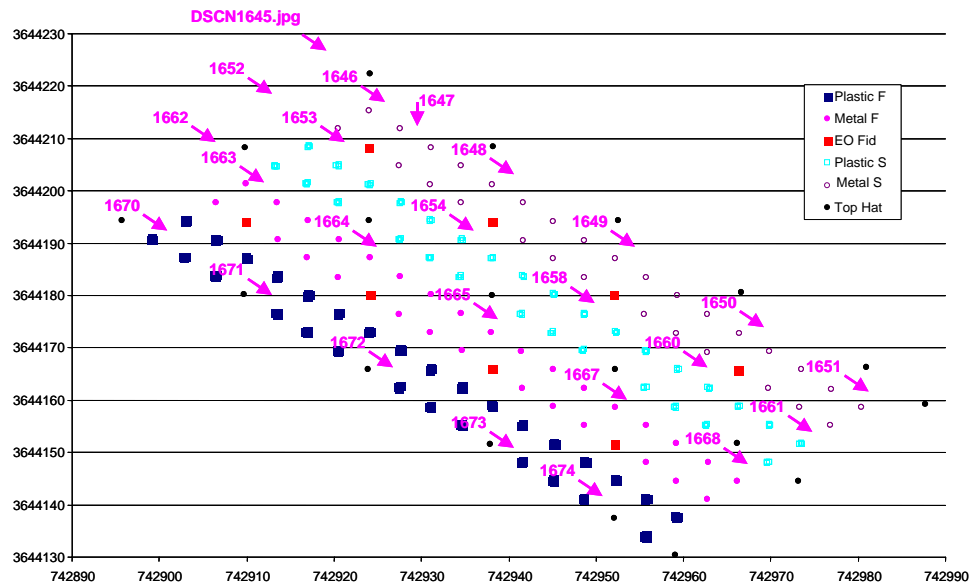


Figure 4-13: Photo capture locations

The catalog of images was used to determine the ground structure make-up and vegetation placement within the scene. Additional, close-up photos of vegetation and targets were also available. These photos allowed the author to determine relative sizes of objects and get a general feel for the terrain of the scene. See Figure 4-14.



Figure 4-14: Metallic mine (L); Desert Bush (R)

Coupling the photo location figure was a survey listing of exact positioning for each surface and buried target within the scene. This listing gave northing, easting, and elevation data (all in meters) for each target, which was invaluable in determining the geometric layout of the scene's target and background objects.

Detailed weather data was collected spanning the length of the collection in 15-minute increments, 24 hours per day (see Table 4-1). The data included a per day listing of time, temperature ($^{\circ}\text{C}$), relative humidity, atmospheric pressure, wind direction, wind speed (meters per second), maximum wind speed, vertical solar irradiance, reflected solar irradiance and amount of precipitation (inches). Also included in the data were soil temperature measurements at six different depths (surface, 5, 15, 25, 35, and 45 cm below surface). Notes or descriptions of how the measurements were obtained were unavailable. It is assumed that the weather information was most likely obtained using a portable weather station and the soil temperature information using contact thermocouples.

U.S. ARMY YUMA PROVING GROUND
DOD COUNTERMINE TESTING AND TRAINING RANGE
METEOROLOGICAL SERVICES

Date: 2003-04-03															
TIME HHMM	TEMP	R/H %	PRESS Mbs	WIND DIR Deg	WIND SPEED Mps	WIND SPEED MAX	SOIL TEMPERATURE DEG C					VERT SOLAR Wm	REFL SOLAR Wm	15 MIN PRECIP In	
	DEG C						SURFACE	5 cm	15 cm	25 cm	35 cm				45 cm
0:15	14.3	34	996.5	282	5.8	8.9	15.1	21.6	24.8	26.3	26.2	25.6	0	0	0.00
0:30	14.1	36	996.5	272	7.6	9.5	14.9	21.3	24.7	26.2	26.2	25.6	0	0	0.00
0:45	13.9	38	996.4	269	7.2	9.4	14.7	21.1	24.5	26.2	26.2	25.6	0	0	0.00
1:00	13.8	39	996.5	277	4.3	7.3	14.5	20.9	24.3	26.1	26.2	25.6	0	0	0.00
1:15	13.5	40	996.3	276	4.4	7.4	14.4	20.7	24.2	26.1	26.1	25.6	0	0	0.00
1:30	13.4	41	996.2	275	5.3	8.3	14.2	20.5	24.0	26.0	26.1	25.6	0	0	0.00
1:45	13.2	41	996.2	283	4.0	7.0	14.0	20.3	23.9	25.9	26.1	25.6	0	0	0.00
2:00	13.0	42	996.1	281	4.6	6.0	13.8	20.1	23.7	25.9	26.1	25.6	0	0	0.00

Table 4-1: Portion of WAAMD weather listing

Radiosonde information collected during the time frame of the experiment was also obtained from the NOAA website, however, it has not been incorporated into final DIRSIG renderings for reasons to be discussed in later sections.

This concludes the description of the ground truth that was collected and/or available to the author. As mentioned previously, it was quite discouraging to discover that emissivity curves and material properties for targets and background objects were not recorded during the experiment or available from outside sources. This presented unforeseen difficulties in building certain aspects of the scene in DIRSIG.

4.2.2 AHI Radiance Images

Details of the AHI portion of the experiment were well documented by the collection team. This provided the author with some ability to derive information that was needed for the simulation directly from the AHI over-flight data. The AHI data is a collection of calibrated radiance images, referenced with a detailed run-list. Each image can be obtained either roll-corrected or non-roll corrected, both calibrated to absolute radiance values in units of 10 times microflicks or $10 \times [\mu\text{W}/\text{cm}^2/\text{sr}/\mu\text{m}]$ (see Figure 4-15). Each data set contains 70 bands of data, ranging from 7.9 to 11.5 microns. The roll-corrected data have also been shortened to view only the area above the minefield being imaged. Raw data was not available directly from the WAAMD FTP site, but this did not pose problems. In addition to the AHI data sets, detailed flight path information determined from GPS/INS data for each set was available as well as an overarching ‘run

list' which detailed cataloging information pertaining to each flight path. The run list includes each filename, file size, number of frames in image, run date and time, test site name, altitude of run, average airspeed, calculated ground sample distance in the along and cross track directions for the run, image width, image height, and any additional comments pertaining to the run. Comments include items such as 'Air Speed Estimated' or 'GPS Error', depending on problems encountered or noticed. After examining descriptions of the available data sets, four data sets were selected and downloaded from the site. They are listed below. All of the downloaded AHI data was acquired on 3 April 2003.

- 700 ft. altitude at 13:11 hrs local time (run #131101)
- 1400 ft. altitude at 13:19 hrs local time (run #131932)
- 700 ft. altitude at 18:55 hrs local time (run #185525)
- 1400 ft. altitude at 19:09 hrs local time (run #190843)

Nomenclature for each DIRSIG rendering which correspond to AHI data sets are the following, respectively:

- n700.cfg
- n1400.cfg
- e700.cfg
- e1400.cfg

These AHI data sets are calibrated, cut over the minefield and roll-corrected. Finally, a briefing by Mr. E. M. Winter, entitled 'AHI at Yuma Proving Grounds' was available on the FTP site, which describes the AHI data and file naming convention. This briefing was invaluable in determining which data sets would be best suited for modeling as well as providing additional information pertaining to AHI configuration.

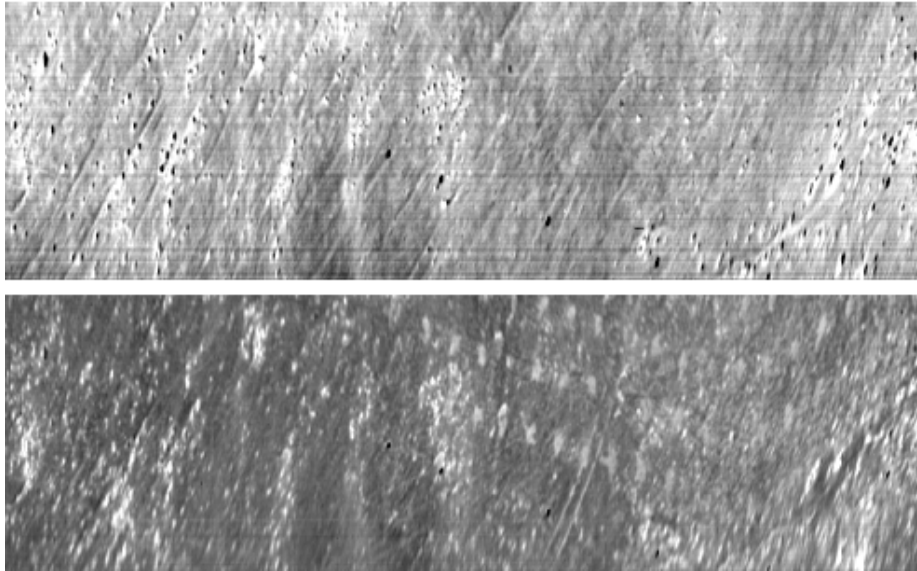


Figure 4-15: Non-roll corrected AHI images, 11.0 μ m (top) and 9.2 μ m (bottom), (Winter, 2003)

4.3 DIRSIG Scene Construction

This section will explain how the DIRSIG scene was constructed and what underlying assumptions have been made to produce final renderings. The synthetic scenes attempt to spatially and spectrally match the truth imagery, however, exact pixel-to-pixel duplication was not the goal. This idea should be kept in the back of the reader's mind.

4.3.1 Terrain and Basic Scene Development

After reviewing the detailed target listing and ground truth photos, it was clear that the scene size needed to be at least 150 by 150 meters to encompass all targets and important background features. The photos showed that topographical variation was extremely minimal over the minefield. In fact, elevation over the entire scene varied by only ± 2 meters, as determined from the target listing. Therefore a basic, flat plate was created with dimensions as listed previously, set at a mean elevation of 0.148 km MSL. Small surface variations were included through the use of a low-scale bump map

(described previously in Chapter 3). The gray-level image used as the base map is shown in Figure 4-16.

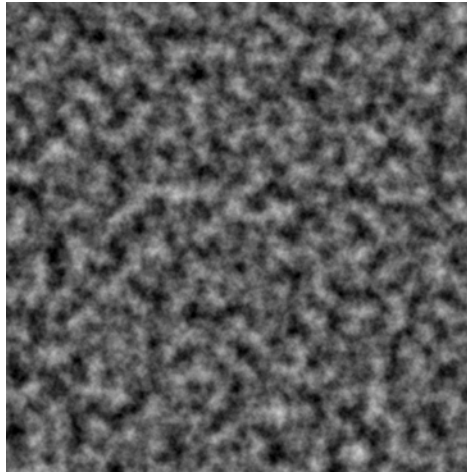


Figure 4-16: Bump map gray level image

The scale value associated with this bump map is set to 0.015, producing a small amount of variation in the deflection of the surface normal vector, ranging between 0° and approximately 17° of total deflection for any given pixel of the mapping image. This is very reasonable based on observations from the ground photos of the truth scene. An additional benefit of adding a bump map to the plate surface is the introduction of thermal variability into the scene due to differences in angular solar loading. For future scene builders, this is a good way of including additional variability into a scene rendered in the LWIR. The mapping image size is 1500 x 1500 pixels with a GSD value of 0.1. This configuration sets the mapping image directly over the scene, without overlap or uncovered space. A mapping image that is too large is not a problem, however a mapping image that is too small causes DIRSIG to “tile” the image to fill the region being mapped. This can cause unwanted tiling artifacts at the edges of the image tiles in final scene renderings. Therefore, if a large area needs to be mapped at a high resolution, it is better to create a very large mapping image that will cover the area rather than use a small map image and allow the tiling process to fill the area. Unfortunately this can be disk space intensive if many different maps are being used in the scene, however, the

results are much more natural. This idea will be revisited in an upcoming discussion on texture mapping.

Latitude and longitude coordinates of the center of the synthetic scene were obtained by converting Northing and Easting coordinates from the detailed target listing into Lat/Long. These values are located in the scene .cfg files. The date and time of each simulation was determined from the AHI run listing corresponding to the run for which each .cfg file is based on. To eliminate some confusion pertaining to the dates listed in the .cfg files, for evening shots, each AHI data set was collected on 3 April 2003; therefore at the time of day the evening shots are rendered it has become 4 April 2003 in the GMT time zone. This requires 4 April 2003 to be the input date for the evening .cfg files.

4.3.2 Target Development and Scene Geometry

A truth listing of target dimensions did not exist, so size estimates were made for both surface mine types, the Top Hat radar reflectors, EO/IR panels and desert bushes. Mine dimensions were estimated from online landmine data sheets (Norwegian People's Aid, 2003) and from data sheets published in the Jane's Mine and Mine Clearance Reference Book (2002–2003). Dimensions of all other targets were estimated from ground truth photos. Approximate dimensions for all targets are listed in Table 4-2.

Object	Dimensions (cm)	Material ID #
Plastic Mine	33 x 33 x 18	23
Metal Mine	37 (dia.) x 15	24
Top Hat	30 (dia. outer base) x 30 (dia. inner base) x 30	22
EO/IR Panel	60 x 60 x 0.2	21
Small Bush	45 (dia.) x 30	31
Large Bush	85 (dia.) x 60	31

Table 4-2: Scene object dimensions

The material ID number for each object correspond to material entries in the scene's material file. The scene's material file will be discussed at length in subsequent sections. Targets were created using 3D modeling software named Rhinoceros™, version 2.0. This program is similar to AutoCad™. If changes are to be made to the dimensions of

objects, either they will need to be re-created in Rhinoceros, or a scale factor (X, Y, or Z) can be applied to individual objects in the scene .odb file.

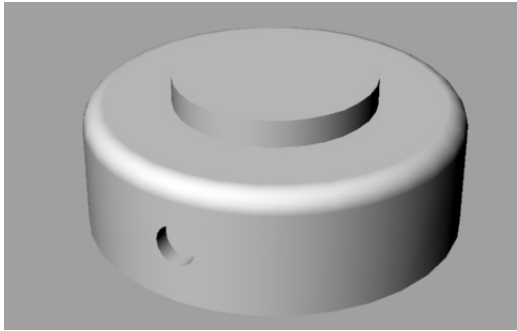


Figure 4-17: Metallic mine model



Figure 4-18: Desert Bush model

Placing individual objects within the scene was accomplished by extracting Northing and Easting coordinates for each object from the detailed ground truth listing, and then converting these into a DIRSIG scene coordinate system. In this new system, the point (0,0) is at the bottom left corner of the scene and the +X direction corresponds to North. Once individual target placements were converted to a common scene coordinate system, the coordinates were inserted into the scene .odb file. The resulting DIRSIG coordinate grid with overlaid target placements is shown in Figure 4-19. Bushes were added by hand using an interactive modeling program, named Bulldozer, after exhaustive examination of the ground truth photos. Bulldozer allows a scene creator to load a base plate or terrain file, then use a point and click method to insert objects. Bulldozer will translate the “clicked” insertion points into the scene coordinate system and list the inserted objects in the scene’s .odb file. There are 165 small and 4 large bushes in the scene, corresponding to locations derived from the ground truth photos. After an initial rendering of the scene was accomplished, it was determined that the initial size of the small bushes was not quite right. In order to make the bushes somewhat larger, a scale factor of 1.5 was applied to each small bush in the scene .odb file. The final dimensions for the small bushes are approximately 45 cm (diameter) x 30 cm (height).

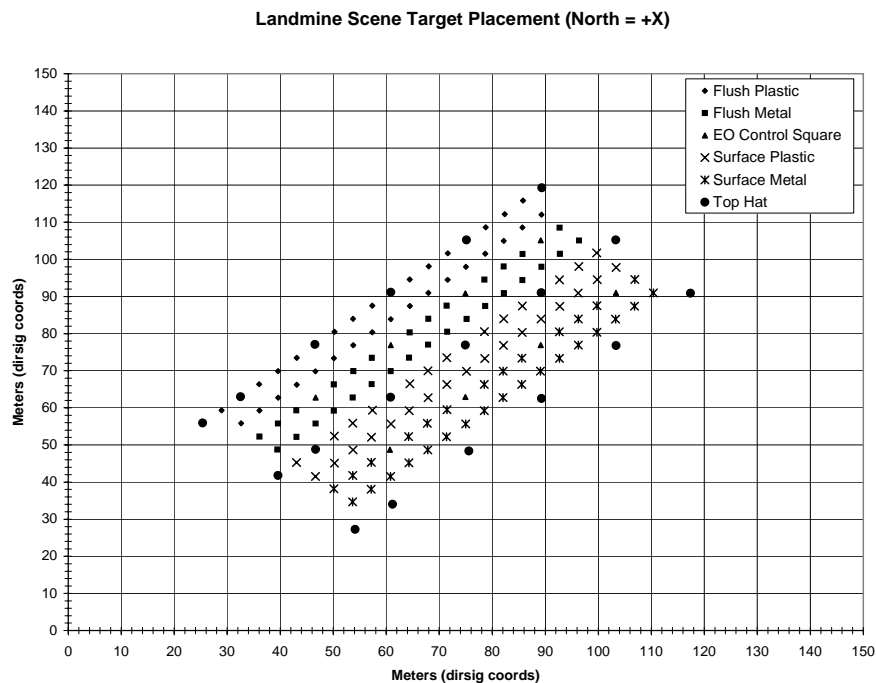


Figure 4-19: Scene target placement and coordinate system

Flush buried mine targets were not added to the scene in the same manner that surface targets were. These targets have been added through the use of a material mapping routine. Rather than attempting to create an area of soil that contains a buried mine using Rhinoceros, material properties corresponding to buried mine phenomenology have been attributed to specific spatial locations on the base plate. The center Northing and Easting coordinates for the buried mine areas were known from the ground truth listing. A gray level, mapping image, 6000 x 6000 pixels, was created that exactly fits over the base plate using a resolution or GSD per mapping image pixel of 0.025 meters. The center point Northing and Easting coordinate for each buried mine area was converted to a specific pixel location within the mapping image. At this location, a small image was inserted, representing the disturbed soil area. This image is shown in Figure 4-20.

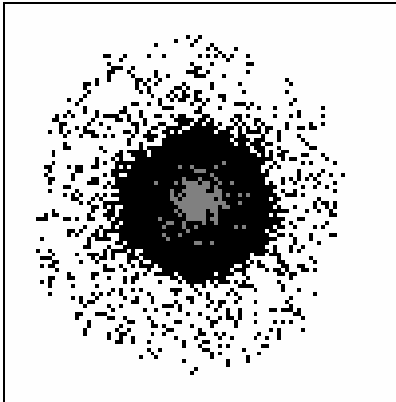


Figure 4-20: Disturbed soil area in material map

The white area corresponds to undisturbed soil, the black to disturbed soil, not directly above the mine, and the gray to disturbed soil directly above the mine. Currently, the gray area is mapped to the exact same material as the black. This region is introduced because it is expected that there are thermal differences between these two areas, however no data has been obtained that allows for differing material parameters to be used that will accurately represent the phenomena. If a temperature map or similar piece of data is developed, it can be easily inserted into the scene. Differences in material parameters will be discussed later. The size of the gray area is approximately 0.4 meters in diameter or slightly smaller once mapped to the base plate. This corresponds to the size of a buried plastic or metallic mine. The distinct black, center circular area is approximately 1.0 meter in diameter, estimated from the AHI data images. As can be seen from Figure 4-20, transition regions have been incorporated to facilitate oversampling and mixing of rendered image pixels. A version of this material map is shown in Figure 4-21 where each dark spot represents the area presented in Figure 4-20.

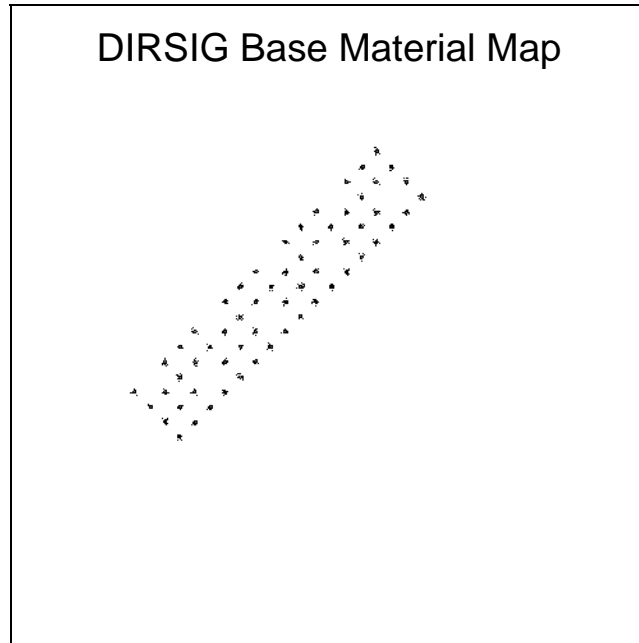


Figure 4-21: Disturbed ground base material map

4.3.3 Weather Data and Atmospheric Characterization

The ground truth weather information listed in Table 4-3 was used to develop the DIRSIG weather file. The ground truth data was converted into a 48-hour weather file using information from 00:00 hrs on 2 April 2003 to 00:00 hrs on 4 April 2003. Only a few conversions were needed to obtain a valid DIRSIG weather file. They include the development of dew point temperatures, the conversion from direct and diffuse solar values from $[W/m^2]$ to $[Langley/hr]$, and the incorporation of sky exposure and cloud type of the scene. The calculation of dew point temperatures is straightforward given the information in the truth table (-1 in the table indicates the calculated dew point temperature was a negative value) and the conversion factor from $[W/m^2]$ to $[Langley/hr]$ is 0.086. Sky exposure was estimated as 1 (full exposure) and cloud type was estimated as 0 (no clouds). Data to the contrary was unavailable. An example section of the DIRSIG weather file with column headings is shown in Table 4-3. The first entry in the

DIRSIG weather file has been taken from the last entry on 1 April 2003 in the ground truth file because the first truth entry for 2 April 2003 starts with time equal to 0.25.

TIME Decimal hrs	TEMP DEG C	PRESS Mbs	R/H fraction	Dew Point	WIND SPEED Mps	DIRCT SOLAR L/hr	DIFFS SOLAR L/hr	Sky Exp	Cloud Type	Precip Type	15 MIN PRECIP cm/hr	Precip Temp
0.00	20.78	989.22	0.29	-1	1.88	0.00	0.00	1	0	0	0	0
0.25	21.08	989.15	0.18	-1	3.55	0.00	0.00	1	0	0	0	0
0.50	20.91	989.07	0.14	-1	3.19	0.00	0.00	1	0	0	0	0
0.75	20.74	989.04	0.13	-1	3.10	0.00	0.00	1	0	0	0	0
1.00	20.32	989.24	0.13	-1	2.14	0.00	0.00	1	0	0	0	0
1.25	19.29	989.33	0.14	-1	0.51	0.00	0.00	1	0	0	0	0
1.50	18.58	989.44	0.16	-1	0.79	0.00	0.00	1	0	0	0	0
1.75	19.03	989.62	0.14	-1	2.20	0.00	0.00	1	0	0	0	0
2.00	18.81	989.69	0.15	-1	2.44	0.00	0.00	1	0	0	0	0

Table 4-3: Portion of the DIRSIG weather file

A modification was made to the original ground truth weather file by reducing the wind speed column by 40%. This was done to make the diurnal temperature prediction of all materials within the scene more realistic. This is a reasonable modification because the thermal model used in DIRSIG predicts surface leaving temperature. The materials within the scene are very close to the ground, if not the ground itself. Also, it is assumed that the weather data was recorded using a portable weather station, which would be monitoring wind speed at a height of three meters according to WMS standards. In a desert environment with high-recorded wind speed, laminar airflow effects would provide a reduced wind speed near the surface, differing significantly from the measured data. The development team decided it was reasonable, given standard material parameters, to adjust the wind speed in an attempt to obtain a more accurate temperature prediction by the thermal model.

The atmosphere for the DIRSIG scene is quite general. A basic desert-like atmosphere has been created and used for the DIRSIG renderings. In the Modtran tape_5 file, the default values have been used other than the following:

- Card 1: used Tropical Atmosphere, Multiple Scattering based at H1, and set the M1 variable (H₂O Profile) to either Sub Arctic Summer or Sub Arctic winter (discussed further below)

- Card 2: used Desert Extinction, Spring-Summer Season, and set the 24-hour average wind speed to 3.343 meters per second (as determined from the weather file)

Setting the M1 variable in Card 1 to an atmospheric profile other than the default variable was decided upon after examination of the resulting radiance curves from initial DIRSIG runs. Compared to extracted radiance curves from the AHI truth imagery, the initial DIRSIG curves appeared highly variable across areas of the spectrum known to contain water vapor absorption features. It seemed quite logical that a baseline “Tropical” atmosphere would contain more water vapor than seen in a typical Arizona atmosphere in early April. At the advice of Dr. Salvaggio, I compared the effect of changing the M1 variable to Sub Arctic Winter, Sub Arctic Summer, Mid-latitude Winter, and Mid-latitude Summer to determine if an alternate profile would align the shape of the DIRSIG generated radiance curves to the AHI radiance curves. After examining the results of the four profiles at the day and evening collection times, a determination was made to use the Sub Arctic Summer water vapor profile for the day renderings and the Sub Arctic Winter water vapor profile for the evening renderings. At the outset, the author hoped to use a single atmospheric profile to sufficiently describe both times of day, but realistically the atmosphere can change quite rapidly and using multiple atmospheric descriptions over the course of the day is closer to truth.

Radiosonde information from NOAA was only available for one time of day during 3 April 2003, which does not allow for atmospheric transition. In addition, there were multiple bad data points in the data. Due to these factors the radiosonde data was not incorporated into the atmospheric profile. An attempt was made to remove the bad data from the radiosonde file, then incorporate the new file into a tape_5 file. Day and evening renderings were generated and evaluated against the truth. Compared to the original “Tropical” atmosphere, the modified radiosonde file did a good job, however these curves were not as good as the generic tape_5 files with the modified M1 variable. Due to the restriction of only one time of day available with radiosonde information and

the large amount of altering needed to get the radiosonde file into the correct format, the author felt that the use of the radiosonde data would be better left out. Had good data been collected at or near the times of day used to render the scene, this data would be excellent to use.

4.3.4 AHI Sensor Parameters

The following section addresses the modeling of the AHI sensor in DIRSIG as well as AHI flight characteristics during the collection and the determination of an estimate of AHI's spectrally correlated noise.

4.3.4.1 Center Channel Wavelength Estimation

The AHI sensor, as stated previously, has the ability to collect 256 bands of data within the 7.5 to 11.7 micron spectral window. From the presentation by Mr. E.M. Winter (2003) it is known that AHI was flown in the 256-band configuration, however, during preprocessing of the data, bands 0 – 9 and bands 220 – 256 were cut out and the remaining 210 bands were binned down to 70. This produces 70 spectral bands ranging from 7.9 to 11.5 microns. Exact center channel wavelengths for the new 70 bands were not provided. In lieu of exact information, approximations to the channel center wavelengths were made. The resulting spectral range from 7.9 to 11.5 microns or a total of $3.6\ \mu\text{m}$ has been divided by the number of channels resulting in a general channel width of $0.05143\ \mu\text{m}$. Using this information, each channel center can be estimated by initially adding half the estimated channel width, namely 0.025715 , to $7.9\ \mu\text{m}$, then incrementally adding 0.05143 to the result, a total of 69 times. The resulting estimates for center channel wavelengths were used to define each of the 70 spectral bands in the DIRSIG response (.rsp) file. The responses of each band were assumed to be Gaussian in shape with a full-width half-max of $0.1\ \mu\text{m}$. The full-width half-max value stems from the resolution information found on the AHI website (2003). The gain term in the DIRSIG response file is currently set to 1×10^7 , which converts the standard DIRSIG

output units of $[W/cm^2/sr/\mu m]$ to 10 times microflicks or $10*[\mu W/cm^2/sr/\mu m]$, the units to which the AHI data were calibrated to. This value can be changed to scale the DIRSIG standard output units to any radiance units the user wishes to have.

4.3.4.2 Ground Spot Size and Inter-pixel Blur Development

To model the pushbroom operation of the sensor and match the compression or elongation of rendered pixels to the truth imagery, additional information was obtained from the AHI website (2003). The focal length of the system is 111 mm, the frame rate is 150 Hz and the IFOV (Instantaneous Field of View) is 0.5 milliradians. This translates to a sensor detector size of approximately $55.5 \mu m \times 55.5 \mu m$ using equation (4-1) where l_d is the length of a side of the detector element and f is the focal length of the system (Schott, 1997).

$$IFOV = \frac{l_d}{f} \quad (4-1)$$

Average aircraft speed and GSD in the along track and cross track directions for each AHI truth acquisition is listed in the ground truth data. Since AHI's frame rate is fixed, the velocity of the aircraft must be tuned to coincide with the flying altitude in order to avoid V/H error, as described in Schott (1997). This did not happen during these collections. At a flying height of 700ft with an IFOV of 0.5 milliradians, a perfectly tuned aircraft speed would produce a GSD in the along track and cross track directions of 0.107 meters, using equation (4-2) (Schott, 1997).

$$GSD = H \cdot IFOV \quad (4-2)$$

In actuality, the evening, 700ft run had an average aircraft speed of approximately 38.125 meters per second. Multiplying by the scan rate, we see that in this case the effective GSD in the along track direction is approximately 0.254 meters. Essentially, the frame rate is not fast enough compared to the aircraft velocity. Projecting the along track and cross track effective GSD's to the surface, it can be seen that each pixel captures a rectangular area. This is illustrated in Figure 4-22 and Figure 4-23.

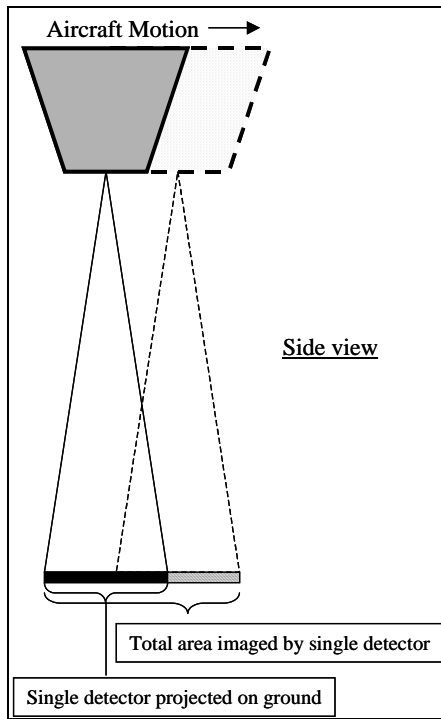


Figure 4-22: V/H Error, Side View

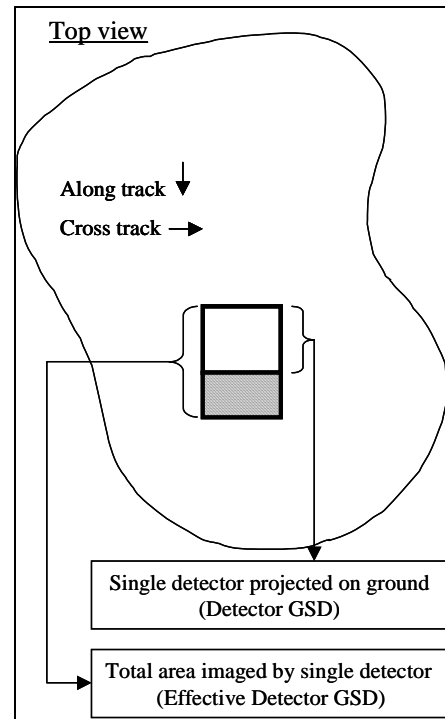


Figure 4-23: V/H Error, Top View

This error results in compression of objects in the along track direction, which is illustrated in Figure 4-24.

Compression effect illustration

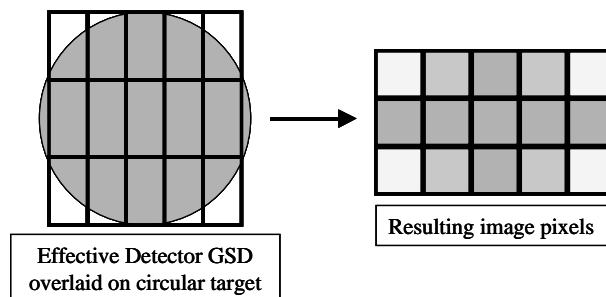


Figure 4-24: Illustration of the compression effect

As can be seen from the figure, a circular target will appear ellipsoidal in a rendered image. This is precisely what is seen in the AHI truth imagery. Each run has a slightly different average aircraft speed, so each run will have a unique effective GSD. To deal with this issue in DIRSIG, the author used the effective GSD in the along track and cross track directions to back-out the would-be length for a side of the detector element. This

detector size was used in the DIRSIG configuration file to create rectangular rather than square detectors. Back to the evening, 700 ft example, the cross track GSD is 0.107 meters, translating to a detector size of 55.5 μm . The along track effective GSD is 0.254 meters, translating to an effective detector size of 132.1 μm . Each of the four configuration files has unique detector sizes corresponding to the truth data's average aircraft speed for that run.

Once the per pixel GSD had been developed, it was necessary to utilize DIRSIG's Oversampling and Point Spread Function utilities to adequately sample this area and produce an accurate representation of the radiance leaving the ground. In a standard DIRSIG rendering, a single ray is cast from the focal plane of the sensor to the center of the to-be rendered pixel, as described in section 3.2. Using only a single ray will cause the radiance leaving the pixel area to be solely determined by the material encountered by that ray, regardless of the number of materials encompassed by the pixel's area. Figure 4-25 illustrates this point.

3x Oversampling Example

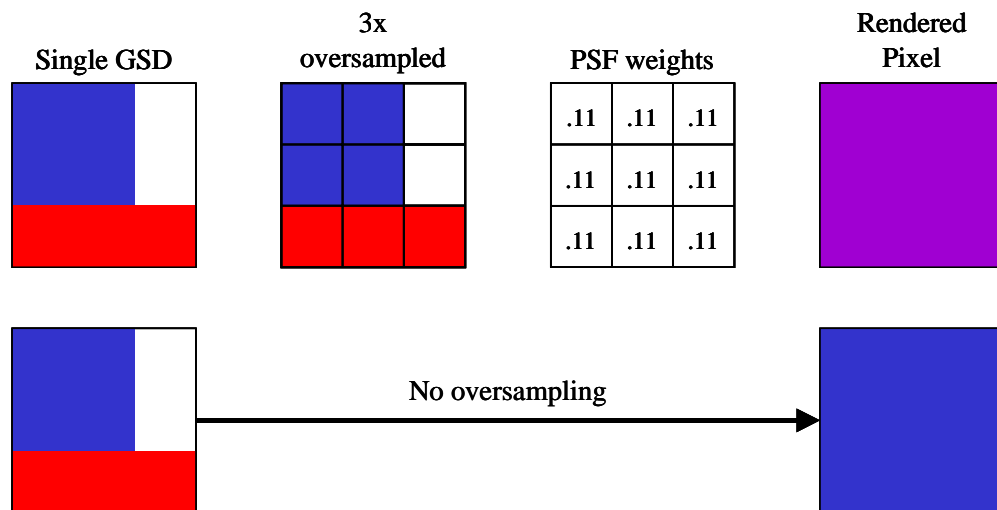


Figure 4-25: 3x oversampling example

As can be seen, without oversampling, the ray that is cast to the center of the pixel will label this area as blue, when in fact it should be purple. The DIRSIG runs in this work

use three-times oversampling (also shown in Figure 4-25), in which nine rays are cast to the pixel area instead of one. A DIRSIG point spread function, or PSF file is used to describe how these nine rays should be weighted to produce the final rendered pixel. In this work each ray is weighted by a factor of $1/9$ or 0.11, producing a uniform mix of the nine rays. The PSF capability essentially “mixed” the rendered pixels.

Pixel mixing allows an accurate representation of the total radiance leaving a pixel area within the scene, however this does not address the issue of inter-pixel blur as seen from the truth imagery. To address this and include it in the final images, AHI’s angular resolution was obtained from the AHI Homepage (2003). This is listed as 0.9 milliradians in the cross track direction and 2.0 milliradians in the along track direction. This is obviously larger than the IFOV of 0.5 milliradians. DIRSIG version 3.5.3 does not facilitate the blurring of rendered pixels with one another, so an alternate method was developed which is applied after the DIRSIG rendering process had ended. This essentially consisted of approximating AHI’s angular resolution as a two-dimensional Gaussian function, and then based on this function a convolution kernel was developed that would run across the imagery and sufficiently blur the data. To develop the kernel, AHI’s angular resolution was first extrapolated to the ground from the flying height of each data set using equation (4-2), and then converted to a common unit of meters. For example, in the 700 ft altitude, evening data set, the angular resolution extrapolates to 0.192 meters across track by 0.427 meters along track. Pictorially, the general form of this extrapolation is shown in Figure 4-26, which assumes a two-dimensional Gaussian spot size.

Inter-pixel blur illustration

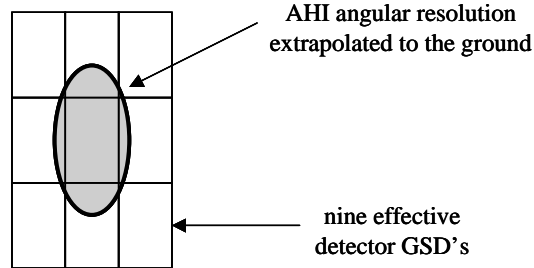


Figure 4-26: Inter-pixel blur illustration

The general form for a normalized two-dimensional Gaussian function is listed in equation (4-3). (Peebles Jr., 2002)

$$f_{x,y} = \frac{\exp\left[-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)\right]}{\sqrt{4\pi^2\sigma_x^2\sigma_y^2}} \quad (4-3)$$

In this case, the variable x represents the across track direction and the variable y represents the along track direction. The extrapolated angular resolution was taken to be a 3σ value in both directions, and then converted into relative amounts of pixels using the effective detector GSD value equal to one pixel in the across and along track directions. This produced a 3σ value in units of pixels. To continue the previous example, the detector GSD is 0.107 meters across track by 0.254 meters along track. Converting the angular resolution in meters to relative amounts of pixels, results in a $3\sigma_x$ value of 1.794 pixels and a $3\sigma_y$ value of 1.681 pixels. These values are used in equation (4-3) to produce a normalized two-dimensional Gaussian function that can be sampled on a 5x5 or 7x7 grid to extract the convolution kernel needed. In this case, both 700 ft altitude kernels were 5x5 in size, while the 1400 ft kernels were 7x7. These kernels have been imported directly into ENVI for use, and were used after the DIRSIG rendering process was complete. The kernels are shown in Table 4-4 through Table 4-7.

Convolution Kernel for Evening 700 ft				
0.000000	0.000200	0.000808	0.000200	0.000000
0.000359	0.023838	0.096494	0.023838	0.000359
0.001770	0.117408	0.475262	0.117408	0.001770
0.000359	0.023838	0.096494	0.023838	0.000359
0.000000	0.000200	0.000808	0.000200	0.000000

Table 4-4: Evening 700 ft convolution kernel

Convolution Kernel for Midday 700 ft				
0.000000	0.000000	0.000195	0.000000	0.000000
0.000273	0.018115	0.073328	0.018115	0.000273
0.001971	0.130715	0.529128	0.130715	0.001971
0.000273	0.018115	0.073328	0.018115	0.000273
0.000000	0.000000	0.000195	0.000000	0.000000

Table 4-5: Midday 700 ft convolution kernel

Convolution Kernel for Evening 1400 ft						
0.000000	0.000000	0.001344	0.005365	0.001344	0.000000	0.000000
0.000000	0.000177	0.011233	0.044840	0.011233	0.000177	0.000000
0.000000	0.000631	0.040153	0.160289	0.040153	0.000631	0.000000
0.000000	0.000965	0.061395	0.245086	0.061395	0.000965	0.000000
0.000000	0.000631	0.040153	0.160289	0.040153	0.000631	0.000000
0.000000	0.000177	0.011233	0.044840	0.011233	0.000177	0.000000
0.000000	0.000000	0.001344	0.005365	0.001344	0.000000	0.000000

Table 4-6: Evening 1400 ft convolution kernel

Convolution Kernel for 1400 Noon						
0.000000	0.000000	0.001480	0.005910	0.001480	0.000000	0.000000
0.000000	0.000182	0.011598	0.046298	0.011598	0.000182	0.000000
0.000000	0.000627	0.039881	0.159203	0.039881	0.000627	0.000000
0.000000	0.000946	0.060196	0.240297	0.060196	0.000946	0.000000
0.000000	0.000627	0.039881	0.159203	0.039881	0.000627	0.000000
0.000000	0.000182	0.011598	0.046298	0.011598	0.000182	0.000000
0.000000	0.000000	0.001480	0.005910	0.001480	0.000000	0.000000

Table 4-7: Midday 1400 ft convolution kernel

4.3.4.3 Spectrally Correlated Noise Inclusion

The inclusion of realistic sensor noise to the DIRSIG imagery was critically important with respect to target detection algorithm performance. Keeping this in mind, attempts were made to acquire some form of truth noise statistics or a dark scan image from the AHI sensor. These attempts failed; therefore actual system noise data could not be used to enhance the DIRSIG data. Without input data to work with, a two-step process was developed which derives an estimated dark scan image directly from the truth radiance images and then uses this dark scan to generate a synthetic noise cube containing the same band to band correlation as the estimated dark scan.

Before thoroughly explaining the noise development process, it is prudent to briefly discuss two important pieces of the process, namely Principle Components (PC) transformations and Minimum Noise Fraction (MNF) transformations, as presented in the ENVI User's Guide (1999). Additional background on both transformations can be referenced in Schott (1997) and Schott (2003).

A Principle Component transformation is used to decorrelate and maximize data variability in multi or hyperspectral data sets. This is accomplished by projecting the original multi-band data onto a new set of orthogonal axes defined by the eigenvectors of the data set. The resulting PC bands of data will be ordered such that the first PC band contains the largest percentage of data variance, the second PC band contains the second largest percentage of data variance, and so on. Schott (1997) points out that latter PC bands in multi or hyperspectral sets tend to contain mostly random noise, therefore using only PC bands that contain a high percentage of the data variability for algorithmic work can greatly reduce computer runtime while eliminating unwanted noise effects. When using the Principle Components transformation in ENVI, it is important to note that the eigenvalues reported are equivalent to the amount of variance in the resulting bands.

The Minimum Noise Fraction transformation, “is essentially two cascaded Principle Component transformations” as described in the ENVI User’s Guide (1999). The first transformation decorrelates and rescales the noise in the data such that the resulting noise data has unit variance and is no longer correlated across spectral bands. The noise covariance matrix is the key piece of information needed to perform this step. This can be acquired from a dark scan, estimated from actual data, or collected by other means. Therefore, in order to accurately rescale the noise during the transform process, accurate noise statistics must be known or be able to be estimated. When practically applying this transformation in ENVI, there exists an option to estimate noise statistics from the data at hand. This process follows the theory that each image pixel contains both a signal and noise component, and that the adjacent pixels contain the same signal component, but a different noise component. The process for developing the value of noise present in each pixel is explained in the ENVI User’s Guide (1999), “A ‘shift difference’ is performed on the data by differencing adjacent pixels to the right and above each pixel and averaging the results to obtain the ‘noise’ value to assign to the pixel being processed.” The ENVI User’s Guide (1999) also points out that the best estimate is generated over a uniform region in the image, where the underlying assumptions are

closer to the truth. The second part of the MNF transformation process is a standard PC transform on the “noise whitened” data. The resulting data space can be thought of as two parts: the first associated with large eigenvalues, containing associated image data, and the second associated with eigenvalues near unity, containing mostly noise. By simply removing the bands of data with eigenvalues near unity and applying an inverse MNF transformation, one can remove unwanted noise while preserving image data and reducing dimensionality.

The two-step process used to develop synthetic correlated noise cubes uses the PC transformation and MNF transformation in tandem to extract noise, rather than eliminate it. As mentioned previously, the first step in developing synthetic noise cubes is to derive a dark scan estimate from the data. A flowchart of the Step 1 process is presented in Figure 4-27.

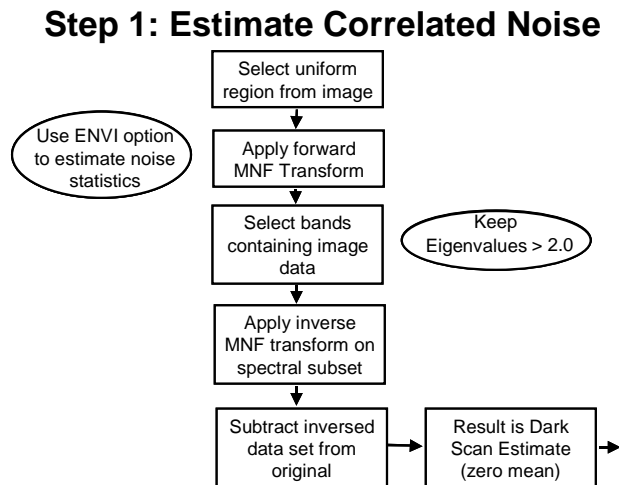


Figure 4-27: AHI noise development flowchart - Step1

First, a uniform region is selected in the image, which happened to be an area of undisturbed soil in this application. Once the region has been selected, an MNF transformation is performed, allowing ENVI to estimate the noise statistics in the data set. From the results, the bands containing eigenvalues less than 2.0 were eliminated, and then the data was inversely transformed to the original data space. The resulting data is

essentially free of noise. To obtain the noise portion of the data, the noise-free data set is subtracted from the original, uniform region data, producing an estimated dark scan. Once a dark scan has been estimated, the process moves to Step 2, as shown in Figure 4-28.

Step 2: Create synthetic correlated noise

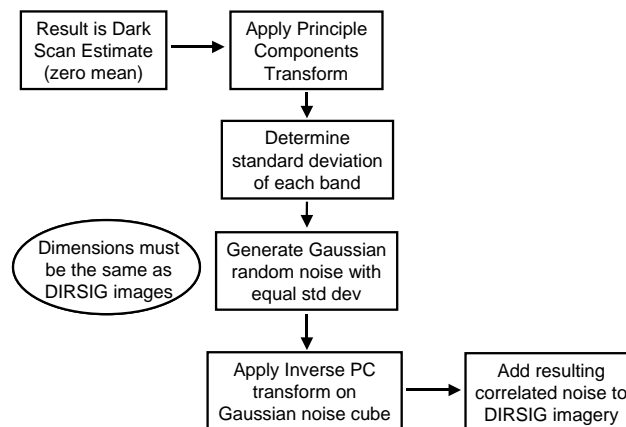


Figure 4-28: AHI noise development flowchart - Step 2

The estimated dark scan is processed through a standard Principle Components transformation to decorrelate the noise. Statistics from the transformed data are extracted to determine the standard deviation in each transformed band. Once the standard deviations are known, a series of Gaussian distributed, random noise images are created containing a zero mean and standard deviation equal to the corresponding transformed band of data. These synthetic noise images are sized to match the spatial dimensions of the rendered DIRSIG image cubes so they may simply be added at the end of the process. The synthetically generated, random noise cube is then inversely PC transformed using the same statistics from the forward transform in order to correlate the synthetic cube equivalently to the estimated dark scan. This process is unique to each truth data set; therefore four synthetic noise cubes have been generated and added to each of the four final DIRSIG image cubes.

In an attempt to determine the accuracy of the noise development process, two comparison evaluations were conducted. The first compares each data set's noise

covariance image from the derived dark scan to the noise covariance image of the synthetic noise cube. These are presented in Figure 4-29 through Figure 4-32.

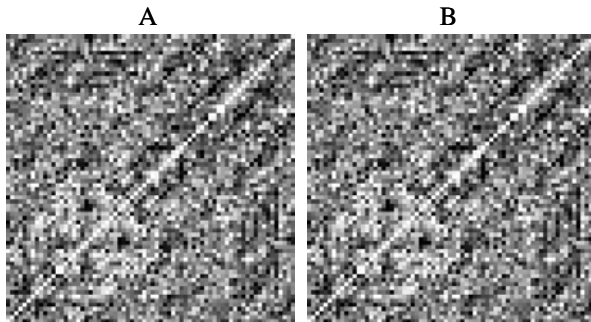


Figure 4-29: (A) E700 AHI derived noise covariance, (B) E700 synthetic noise covariance

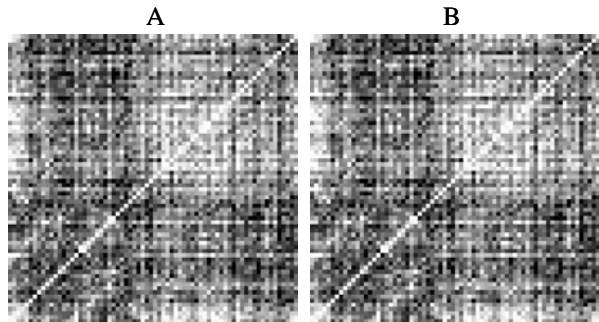


Figure 4-31: (A) N700 AHI derived noise covariance, (B) N700 synthetic noise covariance

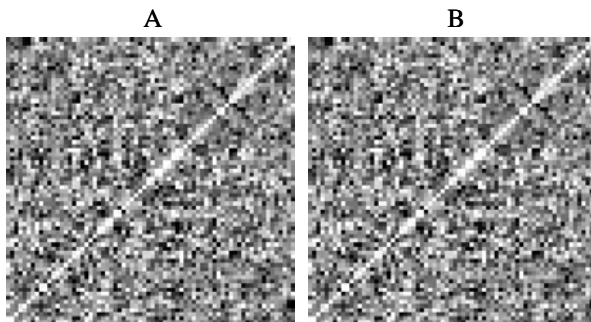


Figure 4-30: (A) E1400 AHI derived noise covariance, (B) E1400 synthetic noise covariance

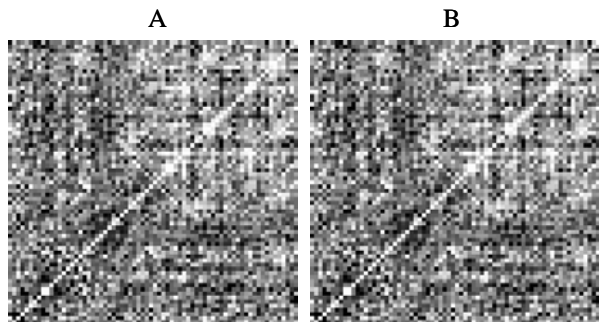


Figure 4-32: (A) N1400 AHI derived noise covariance, (B) N1400 synthetic noise covariance

Clearly, the band-to-band correlation is nearly replicated during the noise generation process. The only drawback to this process is that only proper correlation between the synthetic and derived noise data is ensured. If the derived noise does not contain an accurate representation of the actual noise correlation, the end result will not approximate the truth. In this application, there is nothing that would suggest the derived noise correlation would be different from the actual noise correlation. To ensure that the amount of noise being added is of similar order of magnitude, (i.e. did ENVI's noise estimation process do a good job) plots of signal to noise (SNR) are plotted and compared to a plot of true AHI signal to noise data obtained from Lucey and Winter (1998), where AHI calibration was performed and signal to noise after calibration was

reported. The SNR plots from the synthetic noise data are presented in Figure 4-33 through Figure 4-36 and the truth SNR plot in Figure 4-37.

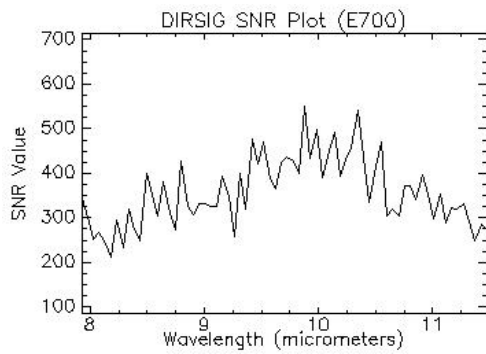


Figure 4-33: E700 Synthetic noise SNR

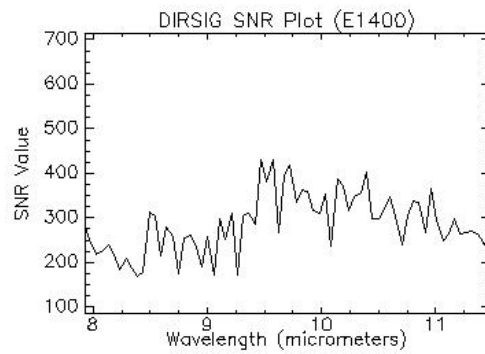


Figure 4-34: E1400 Synthetic noise SNR

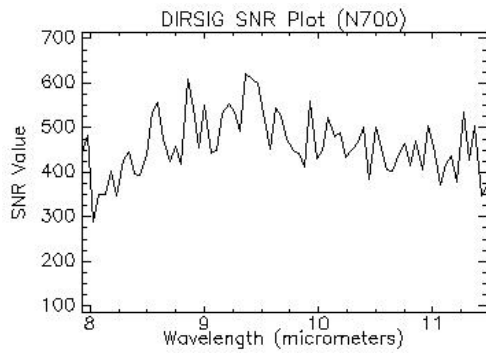


Figure 4-35: N700 Synthetic noise SNR

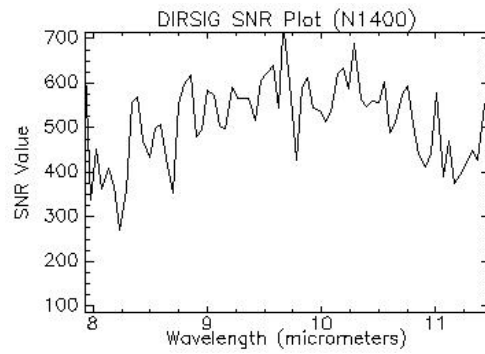


Figure 4-36: N1400 Synthetic noise SNR

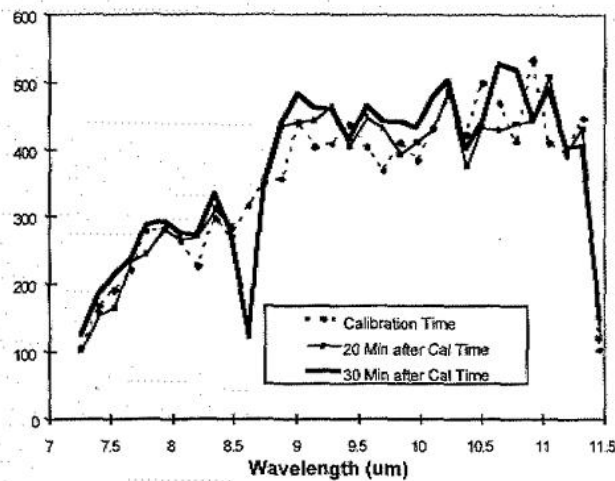


Figure 4-37: AH1 truth SNR plot (Lucey and Winter, 1998)

These plots simply show that the two-step noise estimation process does a reasonable job of estimating an unknown quantity. The SNR values are roughly similar, which achieves the goal for this application. Selecting a more homogenous region for dark scan estimation can improve this process, allowing SNR plots to potentially approach equality with truth data. It is the author's view that in lieu of outside information about the noise of the sensor to be modeled, this method produces a very reasonable result and should be further investigated for use in future synthetic scenes. A drawback to this method of noise estimation is that it does not allow for the incorporation of spatially structured noise. If a full-field blackbody scan was available, the process could be modified to derive the spatial content of the noise, which could then be incorporated into the final DIRSIG renderings.

4.3.4.4 Aircraft Roll Addition and Correction

In addition to straight flight line renderings, DIRSIG has the ability to model aircraft roll through the use of an aircraft profile (.prf) data file. This file lists exact positioning for the imaging platform at specified time intervals in DIRSIG's scene coordinate system. This includes x, y, and z positions as well as roll, pitch, and yaw of the platform. Please reference the DIRSIG User's Manual (2003) for the exact format of .prf files. The .prf file is called in the .cfg file during the rendering process to describe the flight path and motion of the sensor.

Initial DIRSIG renderings used a simple straight-line flight path described by a start point, end point, and number of scans (lines) in the resulting image. These paths were chosen to approximate, not replicate, the flight paths observed in the truth imagery. This produced very reasonable imagery and has been used in the validation analysis. This can be seen in Figure 4-38 and Figure 4-39. Please note that these images may not match scale or contrast exactly when viewed in this document.



Figure 4-38: AHI E700, Band 25

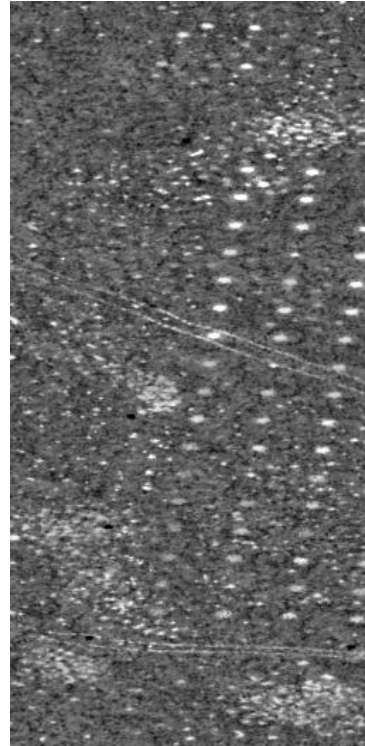


Figure 4-39: DIRSIG E700, Band 25

To add another degree of realism to the DIRSIG data, aircraft profiles were created to add roll to the DIRSIG images, as observed from the AHI images. To accomplish this, a start point and end point were selected to correspond to the general heading of the imaging aircraft. From this, the total distance of the flight path was found then divided by the along track GSD of the run to determine the number of scan lines required in the image. The number of scan lines determined the number of entries in the .prf file. For each entry in the .prf file, a time step was associated, starting with time equal to zero at the start point of the flight path. The truth average aircraft speed for the run was obtained and divided into the total flight distance, resulting in the total time to complete the flight path. The total time was divided by the number of scan lines needed to get the time step between scan lines. This value was incrementally added to the start time, producing a time value for each scan line in the .prf file. A similar method was followed to produce the incremental X and Y values along the flight path. To produce the amount of roll in

the aircraft at each scan line, a few maxima points of roll were determined from the truth imagery along the flight path, then joined with a piecewise linear fit. This estimation process was done by eye and does not exactly replicate the amount of roll in the truth imagery. The piecewise linear fit between maxima points along the flight path produces fractional amounts of pixels of aircraft roll per line, as seen in the truth imagery. To convert amount of pixels of roll to degrees of roll as required by the .prf file format, AHI's IFOV value of 0.5 milliradians was converted to 0.028648 degrees and multiplied by the number of pixels of roll. An interesting note about the process, the DIRSIG User's Manual (2003) defines positive amounts of roll as a clockwise aircraft rotation from nadir as viewed from the rear of the aircraft along the flight line. After viewing the resulting imagery using this convention, the author determined this to be incorrect and exactly opposite of what is needed to produce the expected roll in rendered imagery. Once the .prf files were created, new .cfg files were created which replaced the straight flight paths with a pointer to corresponding .prf files.

Once adequate amounts of roll were simulated, a method for roll correcting the imagery was created. An IDL program, namely "roll_correct.pro", was written which reads in the rendered "rolled" data cube along with a text file listing the amount of pixels each line of the image must be shifted to orient the image appropriately. The program shifts each line of the data cube according to the corresponding text file values and then saves out the new imagery. The text file is pre-generated using the original number of pixels of input roll, rounded to the nearest integer. Residual shift error from the roll correction process exists by design to provide a final result that approximates the precision of correction seen in truth imagery. An example roll-corrected DIRSIG image is presented in Figure 4-40.

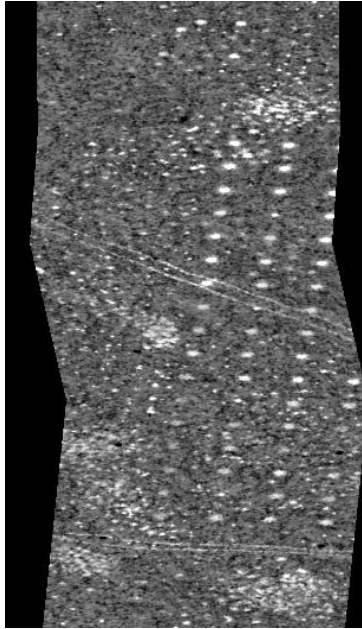


Figure 4-40: DIRSIG roll corrected E700, Band 25

To use the roll corrected renderings for any further processing or analysis, the rolled imagery was blurred using the convolution kernels discussed previously followed by adding synthetic spectrally correlated noise.

4.3.5 Material Parameters

The DIRSIG material file contains all thermodynamic information about a single material needed for DIRSIG's thermal model to develop a material's diurnal temperature prediction. An example entry from the current material file is shown below.

```
MATERIAL_ENTRY {
  NAME = Undisturbed Soil (baseline)
  ID = 30
  SPECIFIC_HEAT = 0.1911
  THERMAL_CONDUCTIVITY = 2.84 # cal/cm/hr/C
  MASS_DENSITY = 1.52 # gm/cm3
  SPECULARITY = 0.0
  THICKNESS = 1.8 # cm
  VISIBLE_EMISSIVITY = 0.76
  THERMAL_EMISSIVITY = 0.91646
  EXPOSED_AREA = 0.368
  EMISSIVITY_FILE = undisturbed_tex_base.ems
  EDITOR_COLOR = 0.8000, 0.6000, 0.3000
  DOUBLE_SIDED = FALSE
}
```


Specifics on how the thermal model interacts with DIRSIG can be referenced in the DIRSIG User's Guide (2003). There is an error in the documentation that is necessary to point out. After exhaustive research, it was learned that the SPECIFIC_HEAT value in the material file must be in units of [calories/gram/degree C] not [Langley's/cm/degree C] as listed in the DIRSIG User's Manual. Also, the THERMAL_CONDUCTIVITY value must be in units of [calories/cm/hour/degree C] and not in [Watts/meter/degree K] as listed in the DIRSIG User's Manual. These changes have been presented to the DIRSIG development staff and should be reflected in future documentation versions. The entry above shows the information that is needed for THERM to provide an accurate prediction of the surface temperature of each facet in the scene. Specific material data was unavailable for every scene material, therefore generic material parameters have been used to represent all materials. These values have been taken from either the suggested values for materials provided in the DCS THERM manual, other published sources, or from the library of materials utilized in other scenes produced and validated by the DIRS group. The following property values have been attributed to materials in the scene as baseline:

- EO/IR Panel = 'paper' as listed in Incropera and DeWitt (1981)
- Top Hat = 'carbon steel – nominal'
- M19 = 'ABS Plastic' as listed at www.arkthermal.com/non-metals.doc
- M20 = 'Carbon Steel – Forest Green'
- Bush = 'Bush – Summer'
- Undisturbed and Disturbed soil = 'Desert Sand'

Broadband thermal emissivity values result from averaging the family of emissivity curves specific to each material over the LWIR region of the spectrum. The emissivity curves generated for the materials in the scene encompass only 7.5 to 11.5 microns, therefore the entire curve was averaged to produce the broadband value. The emissivity curve generation process will be discussed in subsequent sections of this chapter. The visible emissivity parameters were estimated from book values for the baseline material

entries. Once the baseline material file was created, the author interrogated the resulting data sets and refined the material files to produce material temperatures and subsequent radiance curves that fell more closely in line with observed radiance curves seen in the truth imagery. This process was quite iterative and tedious which involved using the thermal model off-line through the “mat-edit” graphical user interface to predict the diurnal temperature curve of each material to be refined, while varying input parameters. The parameters varied were the most questionable input values to THERM. This includes material thickness, visible emissivity, thermal emissivity, and exposed area. Exposed area is a parameter that describes the amount of exposure a given material has with the air. A negative value would represent both sides of a facet exposed, where a positive value indicates only a single side is exposed. Predicted diurnal temperature curves are extremely sensitive to small amounts of variation in this parameter. Since the baseline value was an estimate to begin with, it seemed logical to start by adjusting this value when fine-tuning temperature predictions. The material thickness was also a questionable input value because exact spatial dimensions for objects were unknown. Similarly, the broadband visible and thermal emissivity parameters were questionable. Specific heat, thermal conductivity, and mass density were less unknown, because fundamental material properties are widely known and can be referenced in a variety of sources. Therefore, once the initial material type was known, variations to these parameters were unneeded.

4.3.5.1 Clutter and Thermal Variation Inclusion

After initial evaluation of the scene using the base material mapping scheme, it was determined that additional areas of compacted disturbed soil were needed to facilitate tire tracks as well as add additional thermal variation across the landscape. Due to an unavailability of temperature truth information for the scene, the AHI calibrated radiance images were converted into apparent temperature and apparent emissivity images using the reference channel emissivity calibration utility in ENVI. Band 68 (approximately

11.4 microns) was set equal to a constant emissivity of 0.97, allowing the apparent brightness temperature to be extracted from the data on a pixel-by-pixel basis. Once an estimate of brightness temperature was obtained, the utility could go back and estimate an emissivity value for each pixel based on the assigned temperature. Assigning 0.97 for band 68 was not an arbitrary decision; in fact it was the average emissivity value over the disturbed and undisturbed base emissivity curves in the channel that showed the least variability between the two families of curves. From this data, estimates of the thermal variability within the truth scene were derived and a determination to add additional variability to the scene was made. Plots of apparent temperature versus spatial location in the AHI data are presented alongside initial temperature plots for the DIRSIG renderings in Figure 4-41 through Figure 4-44. It is important to note that the DIRSIG temperature profiles were created from actual surface temperature values as determined by DIRSIG. These values do not take any atmospheric effects into consideration and are not apparent temperatures. The AHI profiles are apparent temperatures as witnessed at the sensor, containing all atmospheric effects. This accounts for most of the noticeable bias observed between the truth and the synthetic. The main purpose for presenting the plots is to illustrate the range of temperature variability in the data.

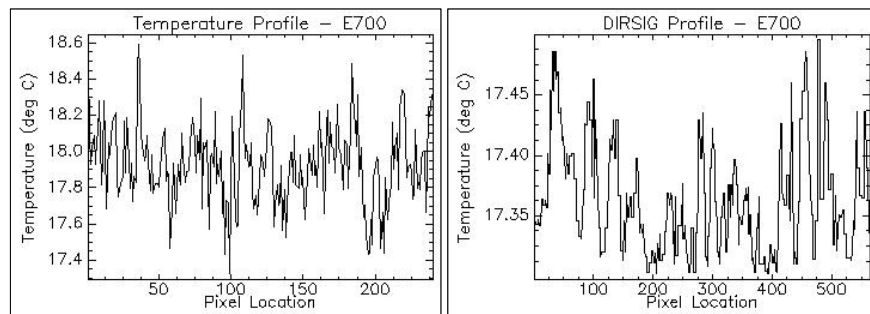


Figure 4-41: (L) AHI temperature profile, E700; (R) Initial DIRSIG temperature profile, E700

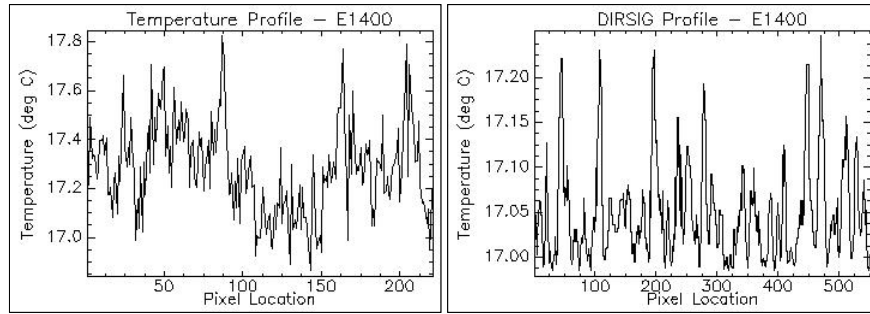


Figure 4-42: (L) AHI temperature profile, E1400; (R) Initial DIRSIG temperature profile, E1400

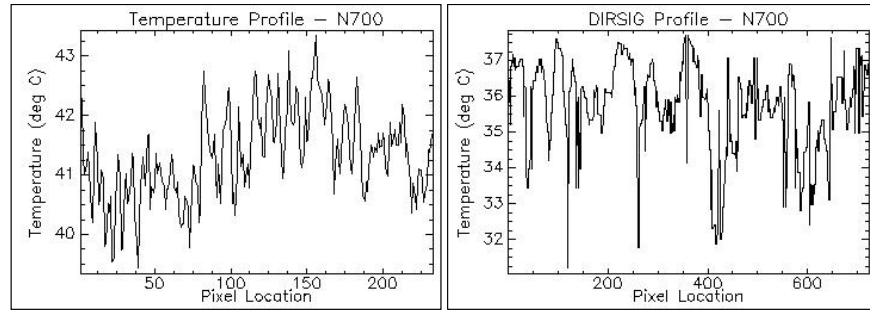


Figure 4-43: (L) AHI temperature profile, N700; (R) Initial DIRSIG temperature profile, N700

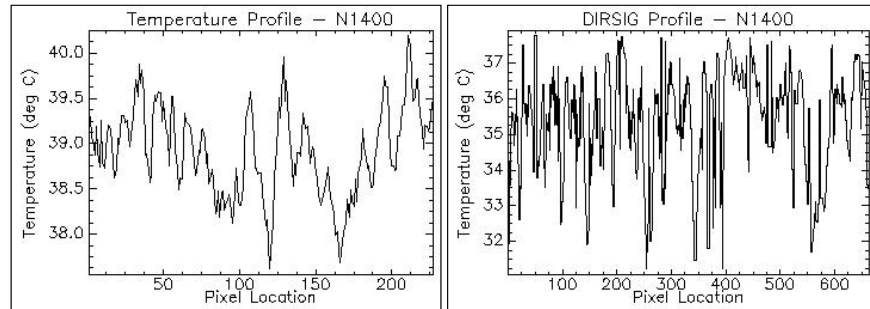


Figure 4-44: (L) AHI temperature profile, N1400; (R) Initial DIRSIG temperature profile, N1400

The evening temperature data shows a significant lack of temperature variability.

Observed variability in the truth data is approximately 0.4°C compared to variability in the synthetic of approximately 0.1°C . The midday synthetic data show a more reasonable temperature variability of approximately 2.5°C compared to variability of approximately 3.0°C in the truth, but is still lacking. This observance in the temperature data combined with a background that also appeared visually less complex led to the development of additional materials, containing variations of their fundamental material parameters.

These varying material parameters allowed the introduction of increased temperature variability and complexity across the scene.

To add areas of disturbed compacted soil, two additional gray level values were added to the base material map (Figure 4-21). The material parameters for the new materials remained equivalent to the area of undisturbed ground with the exception of density values. A medium compact density of 1.75 gm/cm^3 , a highly compact density of 2.0 gm/cm^3 , and the baseline value of 1.50 gm/cm^3 were used to individualize the three materials. This technique provides some degree of thermal variation and adds additional clutter areas in the scene. An example of the resultant material map is shown in Figure 4-45.

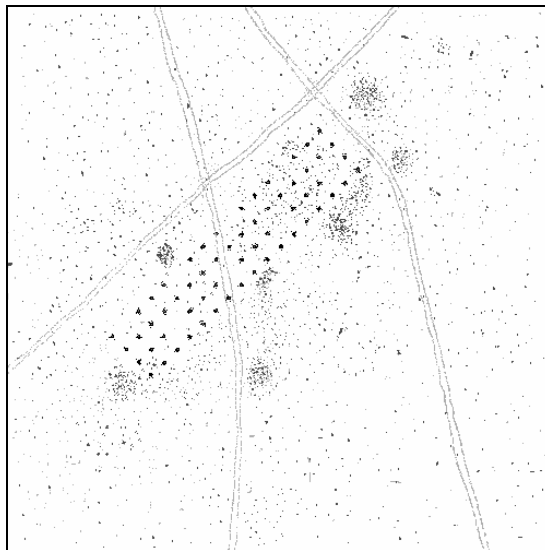


Figure 4-45: Material map with compact disturbed ground areas

After initial creation, this map was slightly modified by reducing the thickness of the “roads” and blending a portion of the disturbed ground areas with the background undisturbed ground, similar to what is observed in the truth data. This can be observed in Figure 4-40, where the disturbed ground areas in the lower left portion of the image blend into the background more than those regions near the top of the image. This added some variation among these regions instead of each being identical throughout the scene. With this mapping scheme, some thermal variation was included, but not enough based on the

apparent temperature data derived from the AHI images. To add additional variability, the undisturbed ground region (all white space in Figure 4-45), was split into three undisturbed ground variants: A, B, and C. The three versions differed only by their input material values for broadband visible emissivity and broadband thermal emissivity. Version B is the baseline undisturbed ground which contains a solar emissivity value (α) of 0.76 and thermal emissivity value (ϵ) of 0.916 (the average value of emissivity curves). Version A is a 'low temperature' variant with $\alpha = 0.73$ and $\epsilon = 0.94$. Version C is a 'high temperature' variant with $\alpha = 0.79$ and $\epsilon = 0.88$. These values were determined through incremental variation of α and ϵ values by realistic amounts until reasonable temperature variation was obtained. The three undisturbed ground materials were mapped to the white space in Figure 4-45, using a version of the bump map image, to which a three-color threshold was applied. The resulting map image is shown in Figure 4-46. Equal amounts of each color are present in the image.

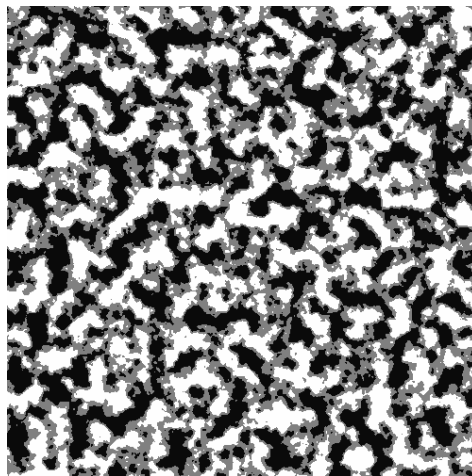


Figure 4-46: Undisturbed ground variant mapping image

The gray represents undisturbed ground variant B, the white, undisturbed ground variant C, and the black, undisturbed ground variant A. A single integrated material map was created by replacing all of the white space in Figure 4-45 with Figure 4-46. This single integrated material map describes the distribution of soils in all current DIRSIG

renderings. The gray level values for the map image with corresponding material ID numbers and descriptions are listed in Table 4-8.

Gray Level Value	Material ID #	Material Description
0	32	Disturbed mine area soil (outer)
119	32	Disturbed mine area soil (above mine)
75	34	Disturbed Soil - Medium compact
175	33	Disturbed Soil - Highly compact
10	35	Undisturbed Soil - Variant A
128	36	Undisturbed Soil - Variant B
255	37	Undisturbed Soil - Variant C

Table 4-8: Material map descriptions

Plots of apparent temperature versus spatial location corresponding to the AHI and final DIRSIG scenes are presented in Figure 4-47 through Figure 4-50. These plots show the progress made in adding thermal variability to the scene.

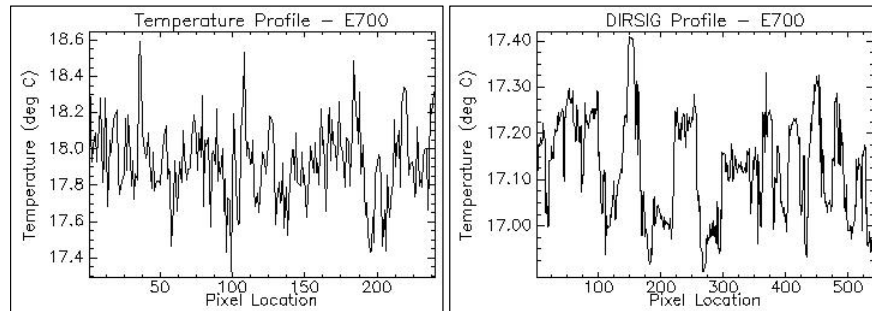


Figure 4-47: (L) AHI temperature profile, E700; (R) Final DIRSIG temperature profile, E700

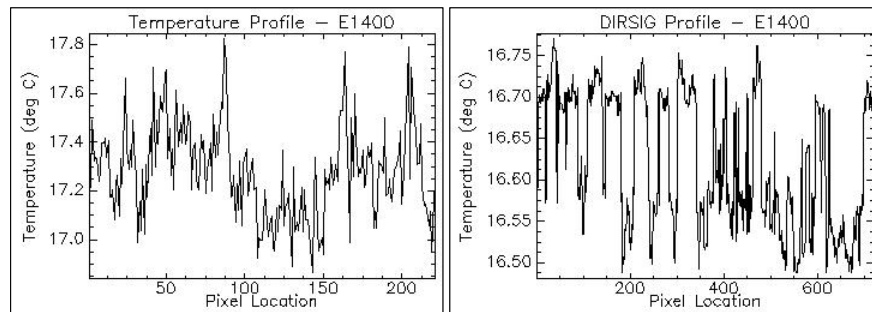


Figure 4-48: (L) AHI temperature profile, E1400; (R) Final DIRSIG temperature profile, E1400

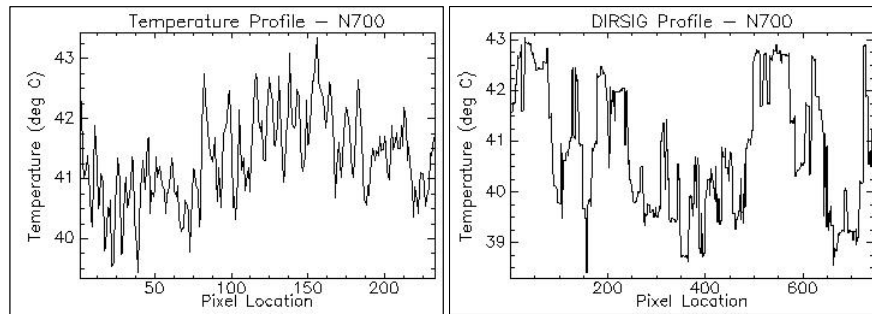


Figure 4-49: (L) AHI temperature profile, N700; (R) Final DIRSIG temperature profile, N700

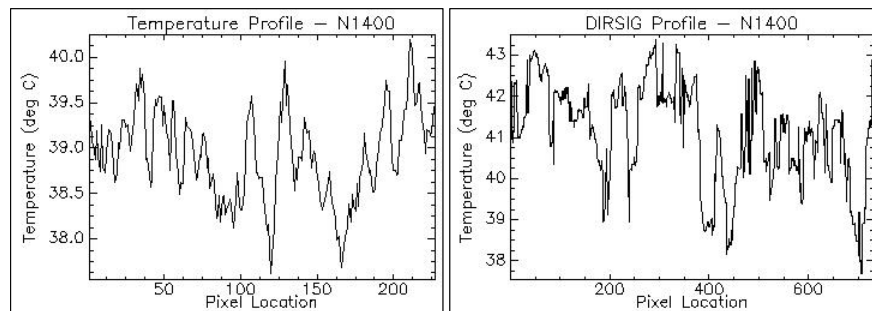


Figure 4-50: (L) AHI temperature profile, N1400; (R) Final DIRSIG temperature profile, N1400

Comparing initial DIRSIG scenes to the final renderings, temperature variability matches much more closely to the truth. As mentioned previously, prior to the creation of new materials, the synthetic scene showed thermal variation of approximately 0.1°C in the evening, 700-foot rendering and approximately 2.5°C in the noon, 700-foot rendering, due only to surface geometric effects introduced through bump mapping. After incorporation of the new material map, thermal variation in the evening, 700-foot rendering was approximately 0.3°C and approximately 3.5°C in the noon, 700-foot rendering. This compares to thermal variation in the truth data of approximately 0.4°C in the evening, 700-foot data and approximately 3.0°C in the noon, 700-foot data. This represents a significant improvement in the thermal modeling process. In addition, the background appears visually more complex, adding to the realism of the synthetic data. The slight bias observed in the final temperature comparison can again be attributed to the truth data being an apparent temperature versus the synthetic being an actual surface temperature.

4.3.6 Emissivity Extraction and Incorporation

Valid emissivity curves for each material within the scene are absolutely essential for successful reproduction of observed material features. In addition, the key buried mine detection feature resides primarily between differences in emissivity curves, so valid curves become even more important. Unfortunately, truth emissivity curves were not available for any of the materials within the scene. Initial thought was to use emissivity curves for disturbed and undisturbed soil collected during the University of Hawaii's Hyperspectral Mine Detection data collections (1996-1998) and attempt to select soil curves that reasonably approximated the soil composition of the truth site. After some investigation, this method would not have provided results good enough for a comprehensive validation of the scene. An alternative solution was developed which takes advantage of an atmospheric calibration technique combined with an emissivity extraction technique, which allowed emissivity curves to be pulled directly from the AHI truth data. The combination of these techniques provides extremely realistic emissivity curves and should be referenced by future DIRSIG scene builders who may not have enough emissivity truth data to work with.

4.3.6.1 Multiple Altitude Atmospheric Calibration Technique

In order to extract emissivity curves from calibrated radiance imagery, an atmospheric correction must first be applied to the data. The truth data for this project was lucky enough to have two data sets imaged over approximately the same area, with a minimal time delay between acquisitions, and acquired at two different altitudes. This situation is needed for a successful multiple altitude calibration, as referenced in Schott (1997). The following discussion follows directly from the treatment in the aforementioned reference. If it is assumed that the collection angle of the imaging platform is approximately straight down, angular correction within data sets can be neglected. In this case, the sensor has a field of view of $\pm 3.65^\circ$ off-nadir and flying heights of 700 feet and 1400 feet AGL, making the preliminary assumption reasonable.

Secondly, it is assumed that contributions due to reflected downwelled radiance are also negligible due to the nature of the imaged scene. Target radiance at the sensor, at a given flight altitude (h) can then be expressed as: (suppressing spectral dependence for clarity purposes)

$$L(h, \theta) = L(h, 0) = \tau(h, 0)L(0, 0) + L_u(h, 0) \quad (4-4)$$

where $L(h, \theta) = L(h, 0)$ due to the small collection angle assumption, $\tau(h, 0)$ is the atmospheric transmission at altitude h , $L_u(h, 0)$ is the upwelled radiance component at altitude h , and $L(0, 0)$ is the extrapolated surface leaving radiance. This also assumes that L_d (downwelled reflected radiance) contributions are negligible. $L(0, 0)$ is determined by using the target radiance value at each flight altitude, then linearly regressing to an altitude of zero. Equation (4-4) can be thought of as a simple line equation in slope-intercept form with atmospheric transmission as the slope and upwelled radiance as the intercept. To solve for these variables, at least two distinct targets must be identified in the data to provide two equations with two unknown values. If we call these targets A and B we can express this system of equations in the following manner:

$$L_A(h) = \tau(h)L_A(0) + L_u(h) \quad (4-5)$$

$$L_B(h) = \tau(h)L_B(0) + L_u(h) \quad (4-6)$$

Where $L_A(h)$ and $L_B(h)$ are observed radiance curves in the truth imagery at a given altitude h , $\tau(h)$ is the atmospheric transmission at altitude h , $L_A(0)$ and $L_B(0)$ are the previously determined ground leaving radiance values for each target, and $L_u(h)$ is the upwelled radiance at altitude h . This system of equations can be solved for the transmission and upwelled radiance values. In this work, the evening, 700-foot data was used to develop an average scene transmission and upwelled radiance because the images were more conducive to selecting similar targets and there were no residual effects from direct solar radiation that may have introduced artifacts into this process. Initially, single pixel targets were attempted for use, however due to differing GSD's between the 700 and 1400-foot imagery it was impossible to select the same pixels in both data sets. To

alleviate this problem, distinct regions were selected then averaged to produce a single target radiance value. Rather than using the minimum two targets, five target regions were selected during this process. With five regions, four different estimates for $\tau(h)$ and $L_u(h)$ were obtained. These estimates were averaged to produce an average scene transmission and average scene upwelled radiance value for each band of data. This data is plotted in Figure 4-51.

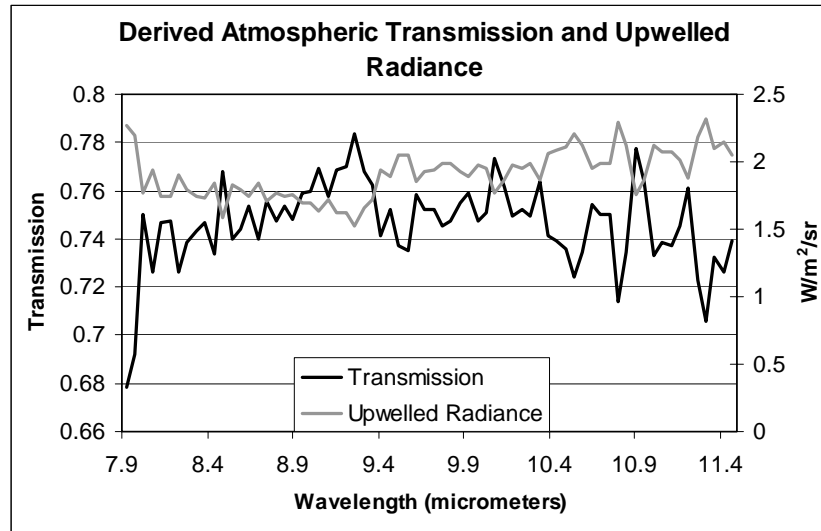


Figure 4-51: Derived atmospheric transmission and upwelled radiance

Average scene transmission and average upwelled radiance can also be derived by linearly regressing the values for observed radiance and ground leaving radiance for each of the five target regions. The slope of the regression becomes average scene transmission and the intercept becomes average upwelled radiance. The former approach tends to average out atmospheric effects more than the latter as additional target regions are used to derive estimates of $\tau(h)$ and $L_u(h)$.

To determine a ground leaving radiance value for any radiance pixel in the image acquired at altitude h , equation (4-7) is applied.

$$L_{Ground} = \frac{L_{Observed} - L_u(h)}{\tau(h)} \quad (4-7)$$

Once an L_{Ground} radiance spectrum for a target of interest is acquired, a Planck Curve Fit is performed to extract the target's emissivity curve.

4.3.6.2 Planck Curve Fitting

The Planck Curve Fit is a temperature/emissivity separation technique discussed in Kahle and Alley (1992) as well as Joseph (1998). This method involves fitting a target's radiance curve to Planck's blackbody radiation curve generated at the highest temperature that will keep the target's emissivity from exceeding unity. This is an iterative process where two initial high and low temperature bounds are selected, and then a blackbody curve corresponding to the mid-point of the temperature range is generated and compared to the target radiance curve. This midpoint temperature then replaces either the upper or lower temperature bound and the process starts over. Iterations continue until the upper and lower bound are within 0.05K (Kahle and Alley, 1992). This method is quick and easy to implement, but hinges on at least one point along the target's emissivity curve approaching unity. Korb et al. (1996) explain that this assumption is valid for most terrestrial surface materials in the 8 to 14 μm spectral range, however many man made materials do not adhere to this rule. In this work, an IDL routine originally written by David Joseph for his MS thesis work (1998), was adapted and used to process atmospherically corrected target radiance curves with the Planck Curve Fit procedure. Figure 4-52 illustrates the procedure for a single emissivity extraction.

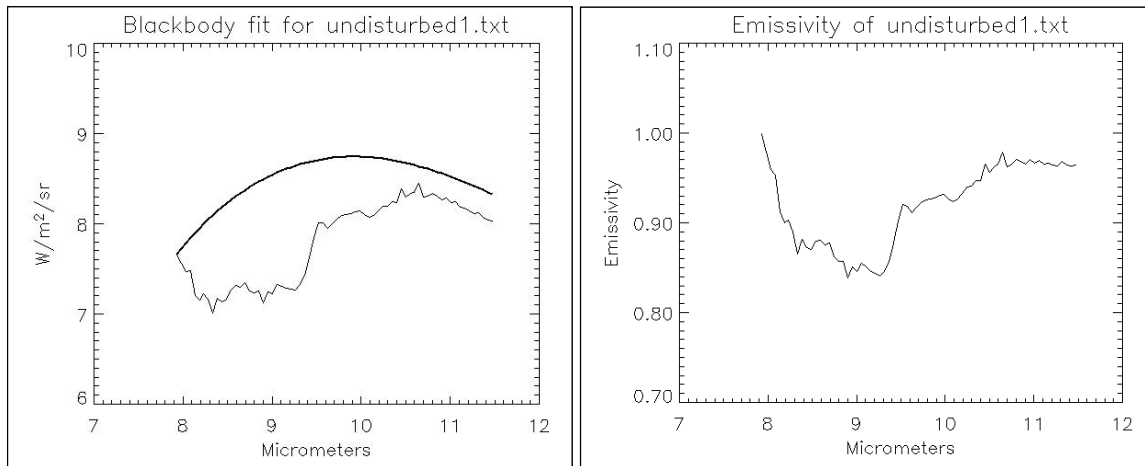


Figure 4-52: (L) Blackbody curve fit for undisturbed soil, (R) Resulting emissivity curve

The most critical spectra to obtain through this process were those of undisturbed and disturbed soil. The true emissivity of both materials does approach unity, so all upfront assumptions were satisfied. This was also true for vegetation spectra, but not for the Top Hat fiducial markers, EO calibration tarps, or both types of surface landmines. With no other avenue for emissivity collection, the Planck Curve Fit was used to derive all spectra used in the synthetic scene. The only exception is the surface metallic landmines, as they could not be sufficiently located in the truth imagery. For these targets, a very general metallic emissivity was applied, obtained from emissivity files generated for other scenes created at RIT. Comparing Figure 4-53 to Figure 4-54, the derived average emissivity curves used for DIRSIG scene creation closely represent the apparent emissivity plots of disturbed and undisturbed soil emissivity as presented by Winter et al. (1996).

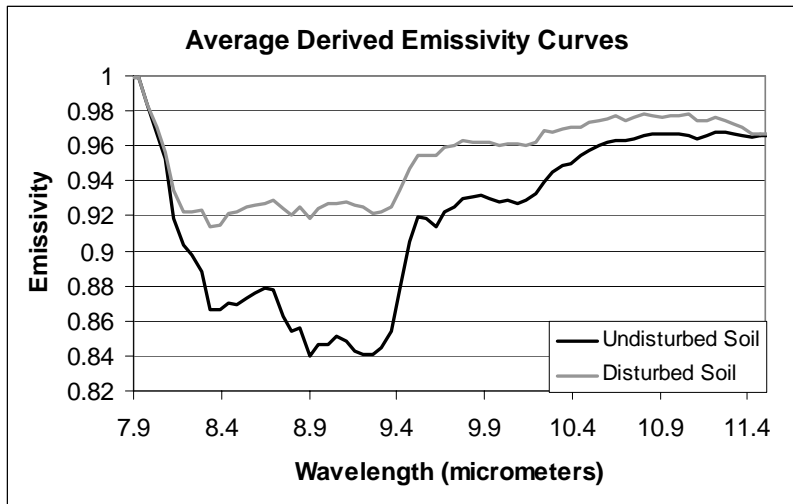


Figure 4-53: Derived emissivity of disturbed and undisturbed soil

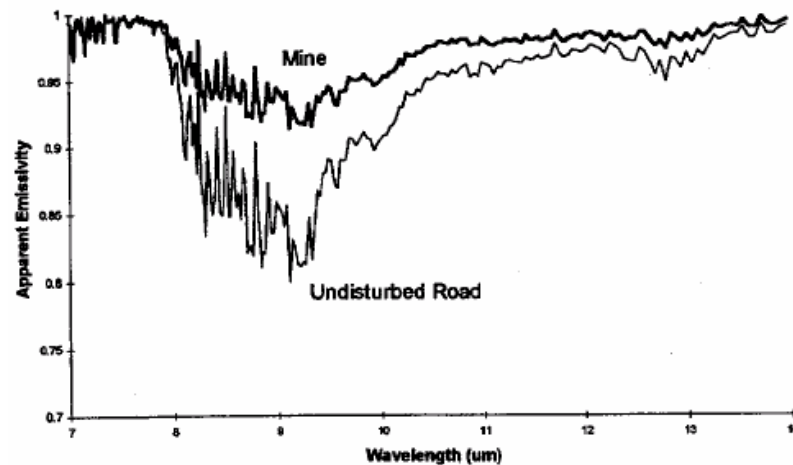


Figure 4-54: Emissivity of undisturbed and disturbed soil (Winter et al., 1996)

4.3.6.3 Emissivity Application

Once the procedure for collecting emissivity curves was finalized, 20 representative emissivity curves were obtained for undisturbed soil and 20 for disturbed soil. These baseline emissivity curves were derived from various regions in the imagery so as to obtain a complete representation of the emissivity variation within each material. In addition to the soil curves, six curves were generated for the Top Hat fiducial markers,

four for the EO/IR calibration panels, six for the surface plastic landmines, and four for vegetation (desert bushes). These baseline families of emissivity curves were then expanded into sets of 500 curves each through the “expand_emissivity_file” utility in the DIRSIG software package. This utility computes multidimensional statistics of the base curve set and then creates a new set of curves based off of these statistics. A formal treatment of this utility is presented in the DIRSIG User’s Manual (2003). The expanded sets of curves were used during the material texturing process (see Section 4.3.7). It is important to note that the emissivity curves generated by this process are only valid over the wavelength region covered by the data sets used to extract them. Therefore, emissivity curves derived in this work are only valid over the 7.9 to 11.5 micrometer range.

4.3.7 Texture

The addition of spectral texture adds emissivity variability across an individual material type by assigning slightly different emissivity curves across material facets. Texturing can also be applied on a sub-facet basis to introduce transition regions between material type boundaries. This process is critical to ensure a statistically realistic representation of the true AHI data. Applying texture to a specific material is accomplished through the use of a gray-level mapping image. The process is quite similar to the mapping process for bump mapping or material mapping (see Section 4.3.5). This process instead maps emissivity curves rather than surface normal deflections or material types. Each gray level has associated with it an emissivity curve pulled from a material’s family of emissivity curves. Each emissivity curve is assigned a gray level based on a Z-score curve selection process over a specific input wavelength range. The details behind texture curve selection using a Z-Score method are discussed in great detail in the DIRSIG User’s Manual (2003), so they will not be presented here. Much difficulty was encountered while deciding on a reasonable gray-level image to use for texturing purposes. Historically, texture has been applied with an overhead image of

the scene to be modeled, but that was not available here. Also, most texture research has concentrated on ways to improve texture capability in the reflective region of the spectrum. In the LWIR, the combination of temperature variation and emissivity variation make the issue even more complicated. The only solution was to use an approximate texture map, see how it looked after rendering and then change maps or make adjustments. The same mapping size issues apply to texture mapping as well as other mapping schemes. This point was brought up in Section 4.3.1, under bump mapping. Originally, the same image used for bump mapping (Figure 4-16), set at a different scale value, was used to texture both disturbed and undisturbed soil regions (all versions). This was because it was the most effective mapping image that was large enough to cover the area, so as to avoid tiling artifacts. After additional research into finding a more suitable texture image, a new “tileable” image (Figure 4-55), as described by the computer graphics community, was obtained which approximated undisturbed soil much better than the original mapping image.

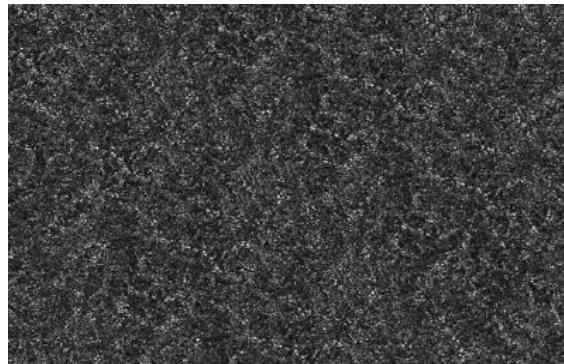


Figure 4-55: Tileable texture image

These tileable images are constructed in such a way that when multiple copies of the image are placed next to each other, the edges blend seamlessly. The use of this new image makes the texture variation for undisturbed soil much more natural looking. It also reduces rendering time because a smaller image can be used, while increasing the scene builder’s flexibility in assigning a GSD for the mapping image. The need to fully cover the area with a single map is eliminated, so the scene builder can use any resolution value

he or she wishes without fear of introducing tiling artifacts from an image that is too small at a small GSD. The new image is used to texture the undisturbed soil regions in the scene. The map uses a GSD of 0.1 meters over a wavelength region from 9.13 to 9.18 μm . This range was selected because it resides at the heart of the Reststrahlen feature and will show the most variability in the family of curves. The old image is still used to texture the disturbed soil regions and Top Hat fiducial markers. The map for disturbed soil regions and Top Hat markers use a GSD of 0.25 meters over the same spectral range as the undisturbed soil map. Tiling artifacts are of no concern here because the GSD is large enough that the 1500 x 1500 mapping image adequately covers the 150 x 150 meter scene area. The rest of the materials in the scene have not been textured due to their extremely small size. Texture would be of little use in these cases.

5.0 Results

As previously mentioned, four data sets were created for final evaluation. The metrics used to evaluate the scene range from the purely qualitative to the purely quantitative. Both methods are necessary for a thorough discussion of the pros and cons of the current scene creation process. This section will address the following methods of evaluation: visual image comparison, Rank Order Correlation (ROC), radiance curve comparison, dimensionality analysis, Gray Level Co-occurrence matrix (GLCM) analysis, Spectral Co-occurrence Matrix analysis (SCM), and R(x) algorithm performance.

5.1 Visual Image Comparison

The first step to evaluating the goodness of fit between the synthetic and truth data is to simply view the rendering and truth imagery side by side. This simplistic qualitative comparison is the best method available for detecting major global differences between the data. Because these data are hyperspectral, only two of the 70 available bands are shown in the comparison, namely bands 25 and 65. These bands lie at approximately 9.16 μm and 11.22 μm respectively, the first within the Reststrahlen feature and the second outside of the feature. The reader is cautioned that the printing process may affect the ability to faithfully match dynamic range between compared images.

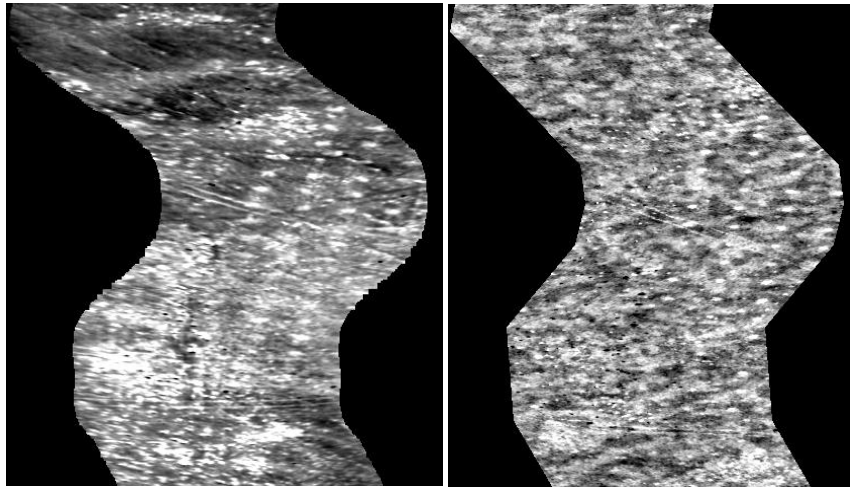


Figure 5-1: Noon at 700 feet, band 25, (L) AHL, (R) DIRSIG

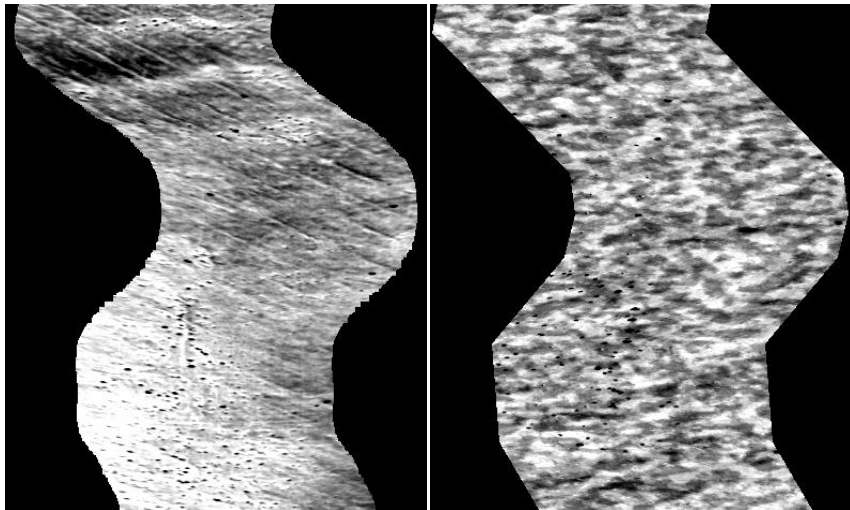


Figure 5-2: Noon at 700 feet, band 65, (L) AHL, (R) DIRSIG

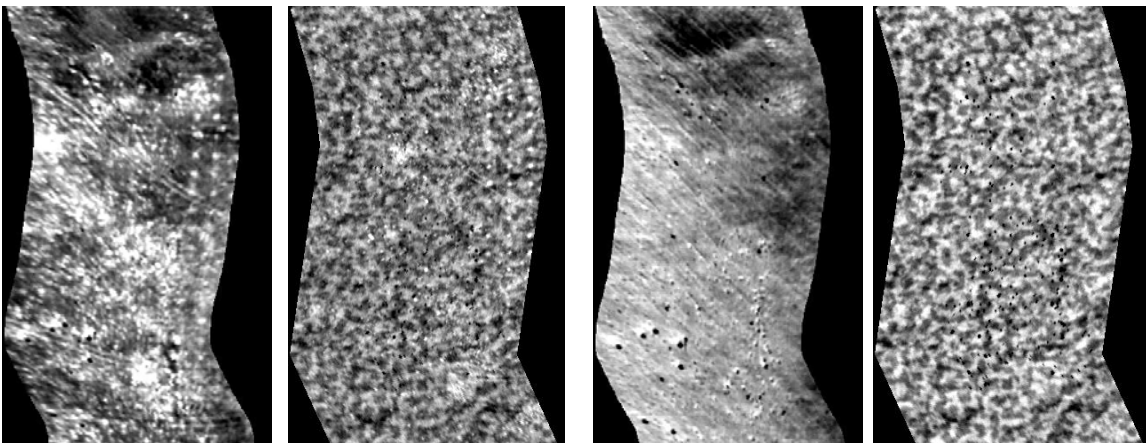


Figure 5-3: Noon at 1400 feet, band 25, (L) AHL, (R) DIRSIG

Figure 5-4: Noon at 1400 feet, band 65, (L) AHL, (R) DIRSIG

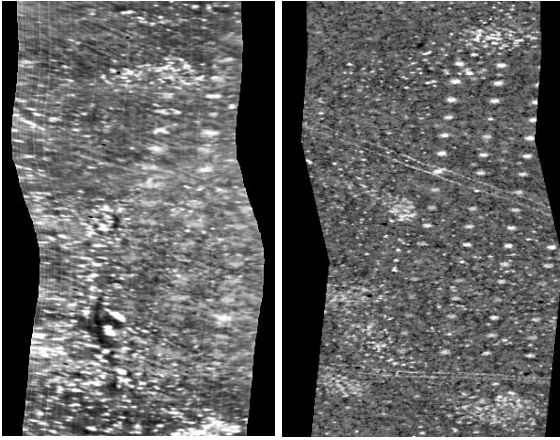


Figure 5-5: Evening at 700 feet, band 25, (L) AHI, (R) DIRSIG

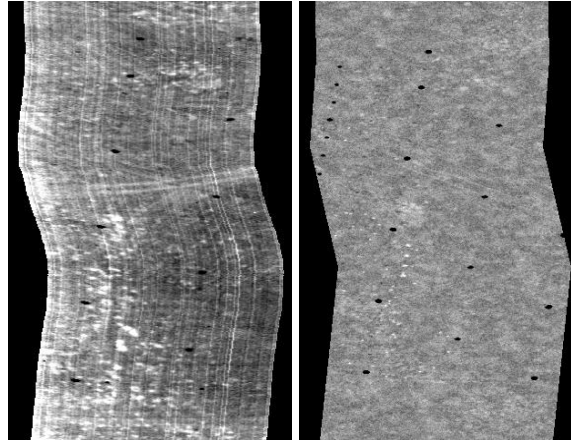


Figure 5-6: Evening at 700 feet, band 65, (L) AHI, (R) DIRSIG

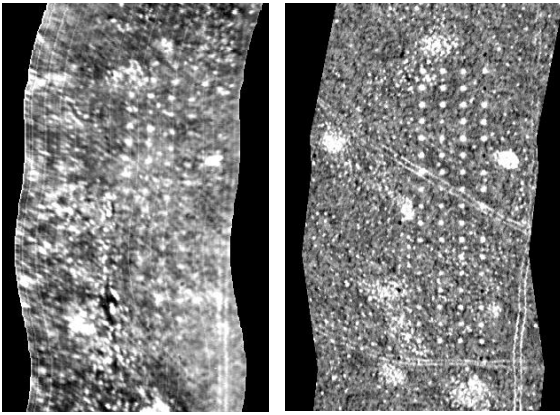


Figure 5-7: Evening at 1400 feet, band 25, (L) AHI, (R) DIRSIG

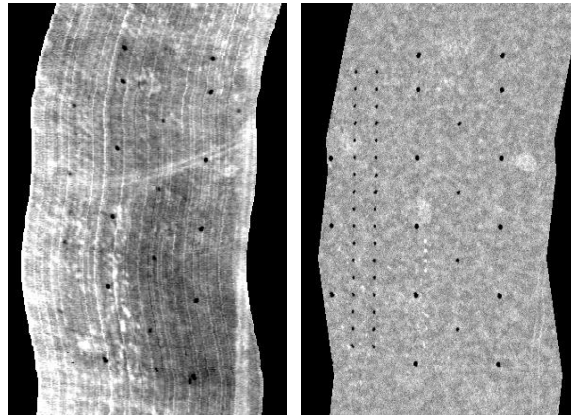


Figure 5-8: Evening at 1400 feet, band 65, (L) AHI, (R) DIRSIG

Figure 5-1 through Figure 5-4 depict the AHI and DIRSIG imagery recorded or rendered at midday, at altitudes of 700 and 1400 feet AGL, while Figure 5-5 through Figure 5-8 depict evening imaging conditions. While comparing images by eye, it is important to show that the images have not been selectively contrast stretched such that they appear to correlate well with each other. To address this issue, histograms of the data presented in images Figure 5-1 and Figure 5-7 are shown in Figure 5-9 and Figure 5-10. These histograms have been clipped such that the zero values (due to the roll correction) were eliminated from the plots. While not every histogram is shown, the presented histograms represent the total data faithfully.

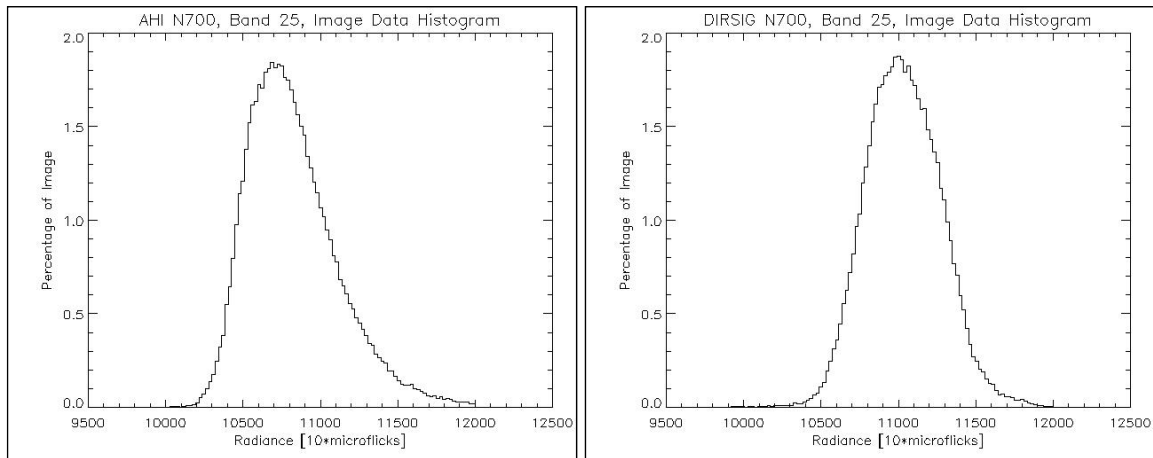


Figure 5-9: Image data histograms from band 25, midday at 700 feet, (L) AHI, (R) DIRSIG

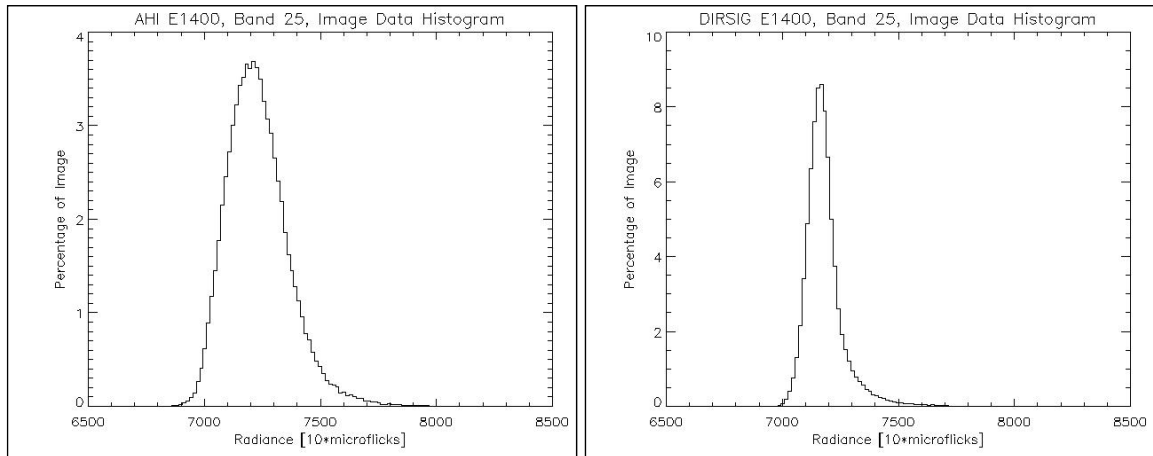


Figure 5-10: Image data histograms from band 25, evening at 1400 feet, (L) AHI, (R) DIRSIG

From a geometric prospective, the DIRSIG data is extremely well matched to the truth. Sizing and positioning of known targets and fiducial markers are practically flawless. Representative flight lines and aircraft roll match well, given that the inputs to this portion of the modeling process were estimates and not intended to be exactly replicated. Additionally, these data show appropriate levels of contrast between target areas (brighter areas in vertical rows) and background areas, including realistic levels of background clutter. Spectrally, it can be seen that the contrast between target areas and the background diminishes when viewing these data off the Reststrahlen feature, as expected. Also, the data show that the process of applying inter-pixel blur has worked extremely well. The target areas in all data sets seem blurred approximately equal to the truth data.

The evening renderings offer good places to observe the emissivity texturing process compared to what is observed in the truth. This is because thermal variability and emissivity variability combine to produce the observed radiance variability in the data. At this particular time of day, there is very low temperature variability across all materials in the scene. Therefore, the observed radiance variability over the background is driven predominantly by emissivity variability. As discussed previously, an overhead image of the truth scene is typically used to drive the scene's texture map. This was unavailable for this work, so a generic texture image was used. The nature of this generality does not allow the rendering to exactly match the truth, nor was this the goal. However, the generic texture image must be reasonable to produce a viable scene. Qualitatively comparing the evening images demonstrates that indeed the generic texture image performs admirably and is a viable modeling option. Additionally in the evening data, the truth imagery shows linear striping noise where the renderings do not. These effects vary spatially with observed spectral band and are produced by inconsistencies between detector elements at the focal plane of the imaging system. Any attempt to characterize this type of noise would require detailed knowledge of the sensor's focal plane at the time of imaging, knowledge that was unavailable for this work. For this reason, structured noise was left out of the simulation. However, as mentioned previously, given a full-field scan of a blackbody with the sensor, spatially structured noise could be derived and included in the scene.

Specific to the daytime data, difficulty was encountered in adequately matching spatial temperature distribution throughout the scene. This is especially noticeable in Figure 5-2 and Figure 5-4, where background brightness variation is spatially on a wider scale in the rendering, resulting from the spatial distribution of soil variants in the material map. As discussed previously, actual variation in temperature is appropriate between the truth and synthetic data, however, the actual spatial distribution of these values is a difficult process to perfect. In order to represent the brightness variation in the background of the rendered scene appropriately, two factors must be considered. The first is spatial distribution of temperature values and the second, spatial distribution of

emissivity values, i.e. the texturing process. When texturing in the reflective region, one only needs to worry about emissivity differences when attempting to spatially match observed brightness variability across a material. In the LWIR, the issue is complicated by temperature variability as well. In these daytime renderings, the material map used to provide temperature distribution is the driving factor for the observed brightness variation. To faithfully produce accurate brightness variability, the spatial distribution of temperature must be reconciled with the texture map, as well as represent the true spatial distribution of temperatures. Otherwise a hot pixel with a low emissivity value may falsely represent a cool pixel with higher emissivity value. The result appears the same, but the phenomenology is incorrectly modeled. The true spatial distribution of temperatures and ultimately variations of soil composition in the scene were completely unknown, requiring a general spatial distribution to be used. In this scene, ground truth was much too limited to fully investigate the issue of texture application in the thermal infrared. Hopefully this work will provide a starting point for further research. That being said, it is the author's opinion that using general mapping images, for texture or material mapping, is an appropriate way to build scenes. As will be described further, the statistical representation of the truth scene is very good. The dilemma is, given a LWIR scene with limited modeling-oriented ground truth, do you build a visually appealing scene, sacrificing the underlying statistical representation, or build a statistically accurate scene at the expense of the way the data appear. The goal is to minimize the error in both directions. Additionally, the ultimate goal of modeling is to create scenes that do not represent an actual "real world" place, but accurately represent physical interactions as would be seen in the real world. So while it is essential to be able to prove the validity of the scene creation process by referencing actual real-world places, an issue with an arbitrary spatial distribution of soil is less of a concern than the underlying phenomenology.

5.2 Rank Order Correlation

One of the most effective evaluations conducted was an evaluation of rank order correlation amongst materials within the scene. The scene is to be processed by

algorithms, keying off likely contrast features in the minefield created by either a thermal difference or emissivity difference. This difference in brightness value is the quantity that is of most importance to correctly model. An attempt to strictly quantify the mean level error between truth and synthetic data will not provide an accurate assessment of how well the synthetic data approximates the truth to an algorithm. As an example, if RMS error were a measure used for evaluation, a constant gain or bias error between truth and synthetic data would yield poor results. If the contrast between scene materials were the same, any global gain or bias would be insignificant to an algorithm. Rank order correlation provides an avenue to evaluate the in-scene contrast while removing error contributions from less significant sources. In order to evaluate the relative brightness of each scene compared to truth, six materials were selected for comparison, depending on the altitude of the data. These materials include vegetation, undisturbed soil, disturbed soil, Top Hat fiducials, EO/IR panels, and surface plastic landmines. The plastic landmines were not large enough to faithfully extract pure pixels of radiance data in the 1400-foot altitude data, so only five materials were used for those data. The number of pixels used for each ranked material in the AHI and DIRSIG data are shown in Table 5-1.

Number of pixels interrogated per material (AHI // DIRSIG)					
	Undisturbed Soil	Disturbed Soil	EO/IR Panel	Vegetation	Top Hat Fiducial
E700	2519 // 8670	570 // 1014	42 // 32	62 // 22	56 // 57
E1400	3766 // 5203	564 // 790	23 // 18	31 // 17	39 // 75
N700	4889 // 9159	425 // 845	27 // 46	24 // 22	34 // 47
N1400	5437 // 9830	184 // 384	21 // 42	24 // 28	25 // 53

Table 5-1: Number of pixels per material used in ROC evaluation

In each band of each scene, these five or six materials were ranked brightest to darkest, then evaluated using the Spearman rank order correlation. The metric values range from -1.0 to 1.0 with 1.0 being perfect correlation, the optimum result for this evaluation. Each band of data carries a unique correlation value, all of which were averaged to produce a single value representative of the entire hyperspectral scene. These results are shown in Table 5-2.

Overall Scene Rank Order Correlation (band average)				
	Mean	Max	Min	Standard Deviation
E700	0.93	1.00	0.43	0.09
E1400	0.94	1.00	0.40	0.09
N700	0.86	1.00	0.60	0.09
N1400	0.93	1.00	0.60	0.10

Table 5-2: Overall scene rank order correlation statistics, comparing DIRSIG to AHI imagery

Results for this portion of the evaluation indicate that brightness contrast between materials can be faithfully reproduced using the DIRSIG model. Minimum values listed for the evening renderings are a singular occurrence, observed in the data's first rendered spectral band, laying on the edge of a water vapor absorption feature. The lowest correlation values are primarily due to the Planck curve fit procedure, as each material's emissivity approached unity in the initial spectral band. This caused all of the derived emissivity curves to be nearly equal in the first band, resulting in a low spread of radiance values between materials. At this time of day where thermal differences are at a low, average radiance values were very tightly grouped, resulting in a relatively arbitrary brightness ranking. Minimum values in the midday data also arise due to a very tight grouping of radiance values for three of the evaluated scene materials, specifically disturbed soil, undisturbed soil, and the EO/IR panels. These three materials tended to be at very similar temperatures at this time of day, leading to a relatively arbitrary brightness ranking in this band. It would be expected that the evening data's minimum ROC values would be lower than the midday data's due to a wider range of average temperatures observed between material types, leading to a less arbitrary brightness ranking. This expectation is supported by the data. The minimum correlation values tend to predict issues with the emissivity derivation process or represent situations where the metric struggles to appropriately rank materials.

In addition to overall scene values, each material was evaluated, ranking the material's radiance value in each spectral band. A rank order correlation value was determined for each material, in each scene and finally averaged for an overall material "score". This evaluation tends to measure spectral correlation determined by atmospheric constituents and emissivity curves. Essentially, this metric evaluates how well the

atmosphere matches the truth atmosphere and how well the derived emissivity curves match the truth. Results for this evaluation are presented in Table 5-3.

Material Rank Order Correlation					
	Undisturbed Soil	Disturbed Soil	EO/IR Panel	Vegetation	Top Hat Fiducial
E700	0.96	0.98	0.91	0.94	0.65
E1400	0.95	0.97	0.87	0.92	0.77
N700	0.85	0.96	0.85	0.96	0.37
N1400	0.86	0.91	0.82	0.94	0.43
AVG Value	0.91	0.95	0.86	0.94	0.55

Table 5-3: Individual material rank order correlation values, comparing DIRSIG to AHI imagery

With the exception of the Top Hat fiducial markers, all individual DIRSIG materials show strong correlation with the truth. Strong correlation across all but one material implies that the atmosphere is modeled quite well. It also points to the Top Hat's derived emissivity curve as the flaw in the process. This was not unexpected, since the Planck curve fit used to separate temperature from emissivity is predicated on the target's emissivity approaching unity, which is typically not the case for man-made materials such as the Top Hat. Taking this one step further, the data implies that the emissivity derivation process tends to break down for non-Lambertian materials. Soils and vegetation, being essentially Lambertian, have high correlation values. The Top Hat fiducials and EO/IR panels are not as diffuse. Particularly, the Top Hat fiducials are quite specular, being made from a polished metal substance. The specularity of the fiducial could increase the contribution of downwelled radiance to the observed radiance curve, causing the atmospheric compensation routine to incorrectly determine ground leaving radiance. This error would manifest itself in the fiducial's derived emissivity curve.

5.3 Radiance Curve Comparison

Another evaluation that is extremely effective for determining spectral correlation between materials is a side-by-side comparison of individual material radiance curves. As with the material-specific rank order correlation, a visual comparison between average radiance curves will clearly show differences in atmospheric composition. In addition, this evaluation method will highlight any global spectral issues associated with the scene. Figure 5-11 through Figure 5-14 are average materials radiance plots spanning five

unique materials, presented with truth data and DIRSIG data side-by-side. These data are identical to the data used for the rank order correlation evaluation. For familiarization purposes, band 1 lies at 7.9 μm and band 70 lies at 11.5 μm .

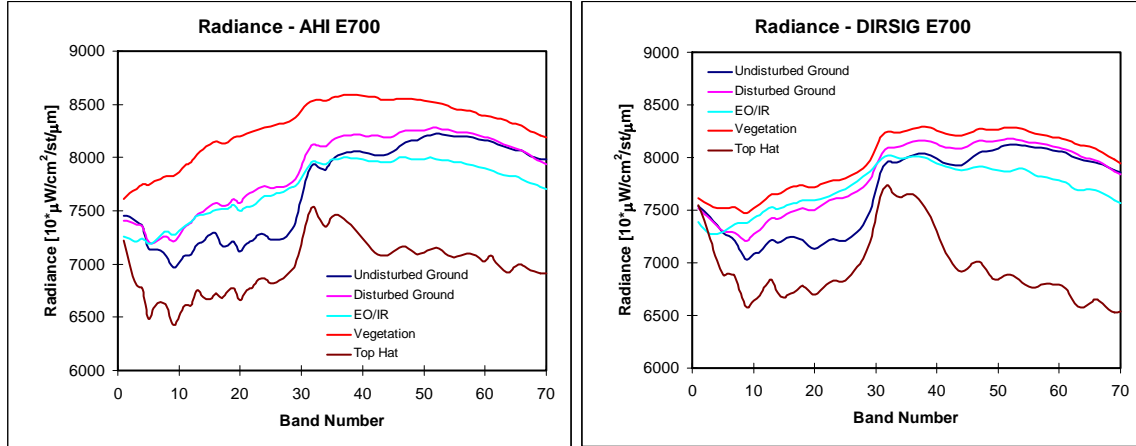


Figure 5-11: Radiance curve comparison of evening data at 700 feet, (L) AHI, (R) DIRSIG

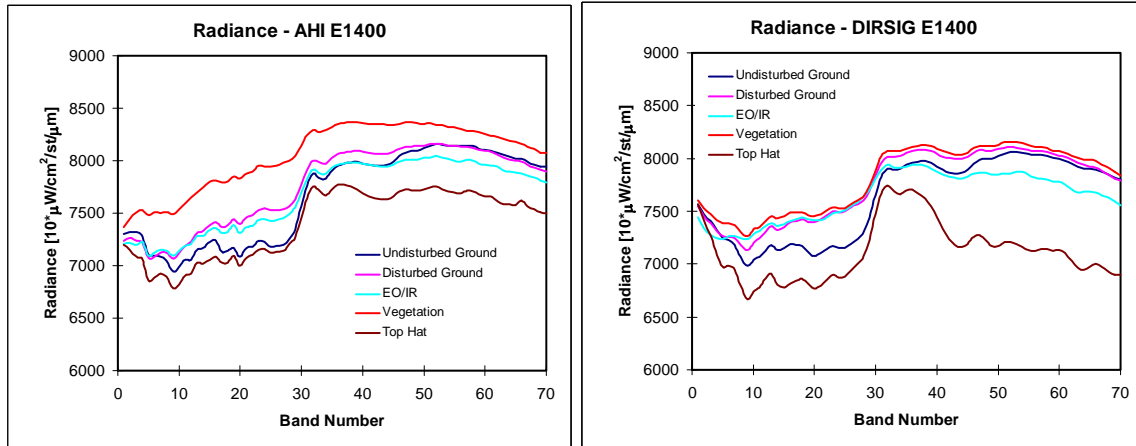


Figure 5-12: Radiance curve comparison of evening data at 1400 feet, (L) AHI, (R) DIRSIG

Beginning with the evening data, characterization of atmospheric constituents appears to be very good. Not every minor absorption feature was intended for replication, so in general, characterization seems quite good over the whole spectrum. Particularly impressive is the near perfect scale in radiance value shown between all truth and synthetic materials. The lack of a global radiance difference also concludes that the atmosphere has been adequately modeled. Vegetation curves and Top Hat fiducial curves do stray from the truth. Vegetation curves show radiance structure that more closely represent disturbed soil rather than a curve that approaches the shape of a

blackbody, as would be expected. This stems from the geometric model of the desert bushes. In the ground truth imagery the desert bushes are very dense, with many tiny branches. These bushes are modeled with a slightly less dense make-up of branches. This leads to radiance emitted from soil surrounding the bushes to mix with the radiance emitted by the bushes themselves. Therefore, pixels that should be pure vegetation are actually mixed with soil. Adding to this, there were definite areas that could be identified in the AHI imagery as vegetation pixels. These regions were large enough to interrogate 10 to 20 pixels each for their radiance curves, which were averaged to produce the plots. The DIRSIG data did not contain areas that were as suitable for interrogating a multitude of pure vegetation pixels. To counteract part of this, additional large bushes were added to increase the number of pure vegetation pixels that were available. This was especially important for the 1400-foot data where the resolution was lower, providing even fewer pixels of pure vegetation. Essentially, this issue worsens when the ability to select a large enough number of pure vegetation pixels is reduced. As mentioned in the previous section, the Top Hat fiducials were difficult to align with the truth. Though not perfect, it can be observed that the curves possess similar base structure, with a main “hump” between bands 30 and 40. Similar pure pixel selection issues exist with the Top Hats as with the vegetation. The Top Hats are quite small compared to the size of the scene, getting only a few pixels per fiducial in the 700-foot data. In this case, it was easier to determine pure pixels in the DIRSIG data, given perfect knowledge of their locations. This leads to radiance mixing with undisturbed soil in the AHI data, readily noticeable in the 1400-foot AHI data, less noticeable in the DIRSIG data. As mentioned previously, most of the Top Hat issues stem from the emissivity curve derivation process. Despite minor glitches, the evening data correlate very strongly.

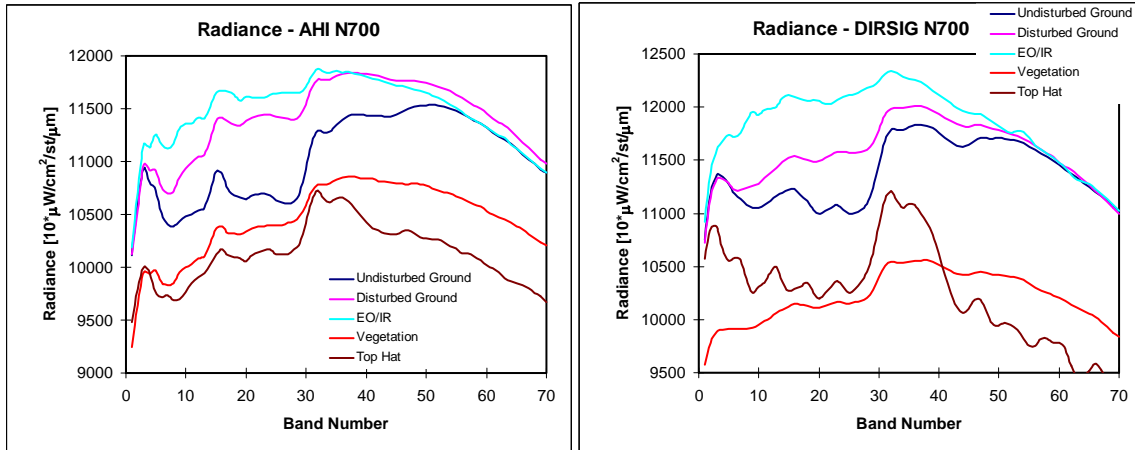


Figure 5-13: Radiance curve comparison of midday data at 700 feet, (L) AHI, (R) DIRSIG

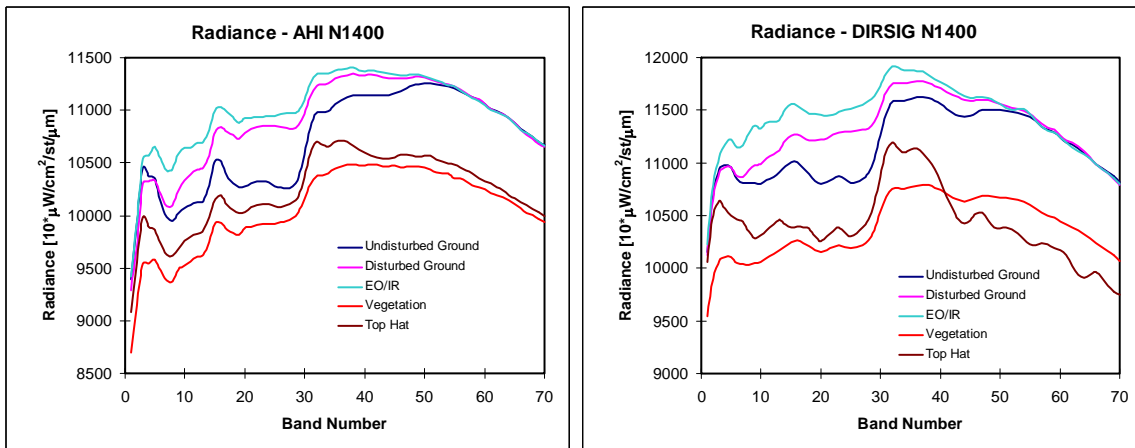


Figure 5-14: Radiance curve comparison of midday data at 1400 feet, (L) AHI, (R) DIRSIG

The midday data correlate quite nicely, but not as strongly as the evening data. There are two global differences that stand out. One, there is a global bias in the DIRSIG data of approximately +500 microflicks and two; spectral features in the data centered about bands 3 (8.03 μm) and 15 (8.65 μm) are more prominent in the truth data. The global bias can be accounted for from a variety of sources. They include slight differences in atmospheric modeling, difficulties in predicting material temperatures at this particular time of day, or weather data misalignment. As discussed in Chapter 4, wind speed values contained in the ground-truthed weather data were lowered by 40% to produce diurnal temperature predictions that were more in line with expected results. This percentage may have been too much of a reduction, producing higher temperatures during the middle of the day. Selecting the percentage of reduction was a trade-off between accurate

temperatures in the evening data with temperatures in the daytime data. This result, while not perfect is the best available. This being said, a global bias is very minor with respect to target detection algorithms and should be taken in context. Secondly, the reduction in major spectral features in the DIRSIG data is due to differences in levels of atmospheric constituents. During experimental trials with alternative atmospheric make-ups, these features were replicated at the same relative strength by lowering the water vapor content in the atmosphere. However, lower water vapor content caused the overall shapes of the radiance curves to be significantly different from the truth. It is more important to accurately characterize the shape of the radiance curve than the strength of the two spectral features. Though diminished, the features are diminished across all materials in the scene. Similar to the global bias, the reduction in spectral feature strength poses little difficulty to a target detection algorithm as long as the reduction is equivalent across scene materials, as observed in this data.

The vegetation curves in the 700-foot data are another point for comparison. The scale of radiance values shows little of the global bias. This essentially falls back to the situation described in evening data results. An average of vegetation pixels in the DIRSIG data emits less radiance than the truth vegetation pixels due to spectral mixing with the surrounding soil. Progressing to the 1400-foot data, vegetation curves are more in line, showing only the global bias as truth vegetation pixels and DIRSIG vegetation pixels are mixed with soil in a similar fashion. A denser geometric structure for desert bushes will eliminate some of the spectral mixing issues.

Top Hat fiducial radiance curves are the most distant from the truth. This issue has been discussed at length and no additional causes of error can be determined from the midday data. It is important to note that while the spectral curves are not completely in line with one another, the general spectral shape has been characterized. In addition, these data are to be used for minefield detection where the fiducial markers are simply intended for orientation within the scene. In all likelihood, these regions would be masked during pre-processing operations or discounted from the detection results. Therefore, knowing it is important to characterize all components of the scene, fiducial

marker correlation is less important than correlation between mined and un-mined regions.

Despite explainable differences in the data, this evaluation overwhelmingly concludes that the DIRSIG scene accurately represents the truth scene from a spectral point of view.

5.4 Dimensionality Analysis

One of the simplest, but most important analyses is an evaluation of dimensionality. The project goal is a good overall statistical correlation between synthetic and truth data. The amount of inherent variability in the data directly compares to the statistical fit between the two. To evaluate this, a standard Principle Components transformation was applied to each of the truth and rendered data sets. If the rendering is a good fit to the truth, the resulting amount of variance in each PC band should be similar between the real and synthetic data. If the synthetic captures most of the data variance in one or two bands, where the truth spreads the variance over six to seven bands, we can conclude that the synthetic is far less statistically complex and will not approximate the scene appropriately to an algorithm. The results of the analysis are shown in Table 5-4, and also graphically as cumulative variance plots in Figure 5-15 through Figure 5-18.

PC Band	Percentage of Variance							
	E700		E1400		N700		N1400	
	AHI	DIRSIG	AHI	DIRSIG	AHI	DIRSIG	AHI	DIRSIG
1	85.80%	79.74%	76.28%	72.99%	84.72%	93.86%	84.53%	96.11%
2	8.64%	14.32%	13.83%	15.75%	13.36%	5.15%	13.38%	3.11%
3	1.12%	1.45%	1.92%	1.28%	0.71%	0.44%	0.66%	0.25%
4	0.53%	0.60%	0.66%	0.75%	0.29%	0.21%	0.35%	0.12%
5	0.33%	0.22%	0.55%	0.68%	0.23%	0.02%	0.22%	0.02%
6	0.16%	0.21%	0.27%	0.64%	0.07%	0.02%	0.06%	0.02%
7	0.14%	0.19%	0.24%	0.49%	0.04%	0.01%	0.05%	0.02%
8	0.12%	0.17%	0.23%	0.46%	0.03%	0.01%	0.03%	0.02%
9	0.12%	0.17%	0.23%	0.44%	0.03%	0.01%	0.03%	0.02%
10	0.11%	0.17%	0.22%	0.41%	0.02%	0.01%	0.02%	0.02%

Table 5-4: A comparison of data variance between AHI and DIRSIG across all sets of data

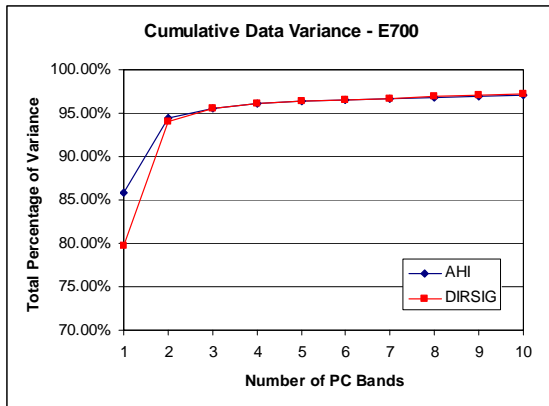


Figure 5-15: Cumulative variance plot - E700

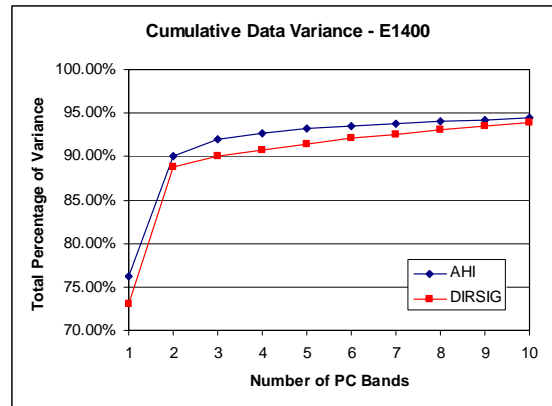


Figure 5-16: Cumulative variance plot - E1400

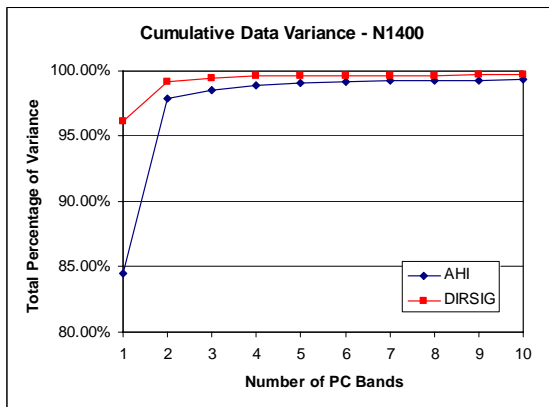


Figure 5-17: Cumulative variance plot - N700

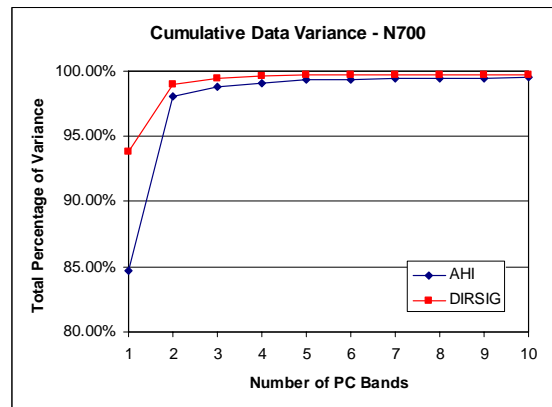


Figure 5-18: Cumulative variance plot - N1400

The evening data compare extremely well, even showing slightly more spread in the synthetic variability. The DIRSIG midday data is not as variable as one would hope. This seems to be due to issues arising with the addition of solar loading in the scene. As mentioned previously, the sun had set at the time of the evening renderings, eliminating solar issues from the data. In the midday data, the added complexity of the sun adds to the thermal variability in the soils. The solution to lacking thermal variability was to increase the number of slightly different soil materials. While this did a reasonable job of approximating the thermal variability, the complexity of the issue may not have been fully replicated. These results point to the need for additional flexibility in the DIRSIG model when adding temperature variation, an issue to be investigated. Additionally, some of the reduction in variability in the noontime data may result from complex surface variation not fully captured by the scene's bump map. In terms of scene building, this could be eliminated with detailed knowledge of the elevation changes across the scene,

incorporated into a Digital Elevation Map or DEM. A DEM combined with a bump map would produce additional surface variability. In this case there was no prior knowledge of the surface variation across the truth scene, so a DEM was unavailable. In general, these results point out issues to be addressed in future iterations of this scene, but are excellent nonetheless.

5.5 GLCM and SCM Analysis

Gray level co-occurrence matrix analysis provides a detailed look at the representation of contrast between disturbed soil and undisturbed soil. The ability to reasonably match levels of contrast between these regions directly relates to an algorithm's ability to distinguish the differences and detect. This evaluation provides a quantitative look at the representation of contrast through statistical comparison between the results. Similar subset regions of the scene were selected for input into the metric. In each scene, a square or rectangular area was selected, containing four disturbed soil patches surrounded by undisturbed soil. These areas were geometrically similar between the truth and synthetic data. Since this texture metric is designed for data that has exact geometric replication, great care was taken in selecting the most spatially equivalent regions. This was not exact given that the synthetic data were not intended to perfectly replicate the truth. The GLCM program operates on a single broadband image; therefore an evaluation was performed on the image from band 26 (approximately 9.21 μm) and the image from band 60 (approximately 10.96 μm) respectively. Band 26 falls in the heart of the Reststrahlen feature, highlighting the contrast between disturbed and undisturbed soil, where band 60 falls in a region of the spectrum where there is little contrast between the two soils. The contrast metric derived from the co-occurrence matrix is essentially comparing target to background contrast between the truth and synthetic data in band 26 and comparing the contrast over the background only in band 60. Each image was processed using the GLCM capability in ENVI. The subset images for band 26 are shown in Figure 5-19 through Figure 5-22, while the subsets for band 60 are shown in Figure 5-23 through Figure 5-26.

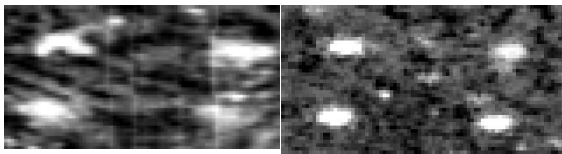


Figure 5-19: E700, band 26 GLCM subset, (L) AHI, (R) DIRSIG

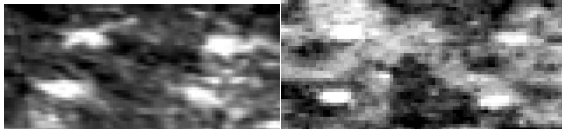


Figure 5-21: N700, band 26 GLCM subset, (L) AHI, (R) DIRSIG

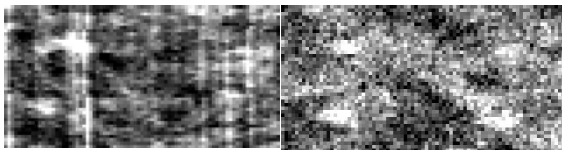


Figure 5-23: E700, band 60 GLCM subset, (L) AHI, (R) DIRSIG

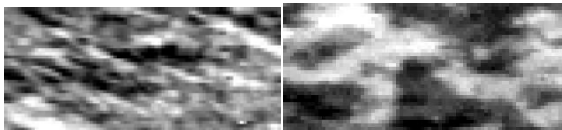


Figure 5-25: N700, band 60 GLCM subset, (L) AHI, (R) DIRSIG

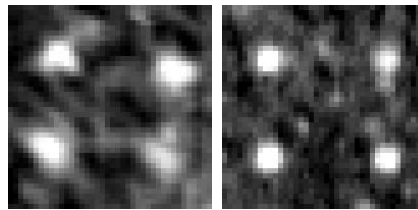


Figure 5-20: E1400, band 26 GLCM subset, (L) AHI, (R) DIRSIG

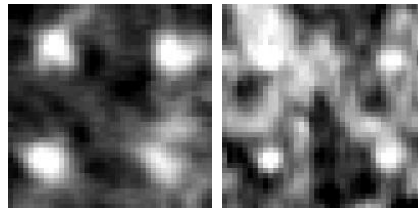


Figure 5-22: N1400, band 26 GLCM subset, (L) AHI, (R) DIRSIG

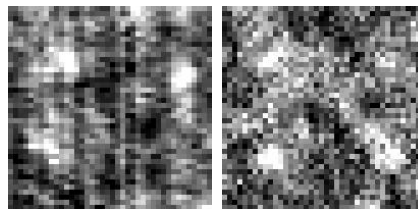


Figure 5-24: E1400, band 60 GLCM subset, (L) AHI, (R) DIRSIG

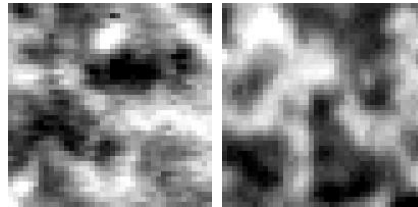


Figure 5-26: N1400, band 60 GLCM subset, (L) AHI, (R) DIRSIG

A 3x3 processing window with co-occurrence shift of $x=1$, $y=1$, as well as quantization level set to 64 were initially used as inputs. These settings are the ENVI default values. The contrast metric derived from the co-occurrence matrix was used to obtain the statistical results. The results from the analysis of the band 26 data are shown in Table 5-5 and Table 5-6.

GLCM Contrast Statistics - AHI Data (Band 26)				
	Min	Max	Mean	Standard Deviation
E700	0.000	374.444	42.340	49.656
E1400	0.000	310.667	24.142	36.743
N700	0.000	359.556	34.935	48.243
N1400	0.000	229.000	23.682	35.756

Table 5-5: GLCM contrast statistics from band 26, AHI data

GLCM Contrast Statistics - DIRSIG Data (Band 26)				
	Min	Max	Mean	Standard Deviation
E700	0.000	444.222	32.995	45.052
E1400	0.000	349.333	27.477	48.799
N700	0.000	342.444	38.103	39.827
N1400	0.000	230.667	24.913	28.370

Table 5-6: GLCM contrast statistics from band 26, DIRSIG data

Based on the mean value statistics generated from the contrast metric, these sets of data fall in line with one another. There are differences in the data, however a good sense of similarity in the results is all that can be expected since these scenes are not spatially equivalent, as the metric would normally require. This evaluation in band 26 shows that the levels of target to background contrast across the subset region selected are very well represented in the synthetic data. In addition to good similarity in general, the data point out that the target regions in the evening, 700-foot, DIRSIG data show less contrast than the truth data. Based on this evaluation alone, these results would predict that an algorithm might have a more difficult time detecting target regions in the DIRSIG data.

Contrast levels in the background were evaluated using the GLCM contrast metric over the subset region from the data in band 60. Initially, the same input parameters were used in ENVI. The statistical results are shown in Table 5-7 and Table 5-8.

GLCM Contrast Statistics - AHI Data (Band 60)				
	Min	Max	Mean	Standard Deviation
E700	0.000	325.333	54.509	47.476
E1400	0.000	346.333	55.091	46.284
N700	0.000	378.889	40.777	39.300
N1400	0.000	315.111	35.478	38.703

Table 5-7: GLCM contrast statistics from band 60, AHI data

GLCM Contrast Statistics - DIRSIG Data (Band 60)				
	Min	Max	Mean	Standard Deviation
E700	0.000	395.444	80.683	52.862
E1400	0.000	439.889	97.251	67.764
N700	0.000	215.444	34.713	35.030
N1400	0.000	145.111	22.023	24.808

Table 5-8: GLCM contrast statistics from band 60, DIRSIG data

Observing the mean values for the evening data, the results show that the levels of background contrast are much greater in the DIRSIG data. The reason behind this is first explained by looking at Figure 5-23 and Figure 5-24. The apparent level of noise in the DIRSIG images is much higher than the level of noise in the AHI images. To confirm this, the signal-to-noise plots generated of the DIRSIG synthetic noise were compared to the historical AHI signal-to-noise plot. These are shown in Figure 4-33, Figure 4-34 and

Figure 4-37. The signal-to-noise is approximately 425 counts in the AHI data at band 60, as opposed to approximately 300 to 325 counts in the DIRSIG data at band 60. The increased level of noise in the DIRSIG data causes the metric to detect a higher level of contrast. This points to the need for improved noise data, especially in the longer wavelength bands. The midday data exhibit the exact opposite trend in band 60. The AHI data show more contrast over the background than the DIRSIG data when comparing the mean values. The lack of contrast in the DIRSIG data points again to the lack of variability in the background. This relates directly to the characterization of the thermal variability in the synthetic data, which has been discussed in Section 5.1 and Section 5.4.

In an attempt to characterize the levels of contrast using a larger and more “target-like” processing window for each subset region, the GLCM evaluation was performed a second time. The input parameters for the 700-foot evaluations used a 7x3 processing window with co-occurrence shift of x=0, y=1, and quantization level = 64. The 1400-foot evaluation used a 5x5 processing window with co-occurrence shift of x=0, y=1, and quantization level = 64. These processing windows approximate the size of the disturbed soil targets in the subset images. The results from band 26 are shown in Table 5-9 and Table 5-10.

GLCM Contrast Statistics - AHI Data (large window, Band 26)				
	Min	Max	Mean	Standard Deviation
E700	0.000	264.000	36.860	35.316
E1400	0.000	95.560	15.153	17.659
N700	0.000	192.905	30.575	34.095
N1400	0.000	83.160	12.459	14.530

Table 5-9: GLCM contrast statistics with large window from band 26, AHI data

GLCM Contrast Statistics - DIRSIG Data (large window, Band 26)				
	Min	Max	Mean	Standard Deviation
E700	0.000	191.952	25.217	29.366
E1400	0.000	123.320	12.942	18.510
N700	0.000	194.905	30.570	29.169
N1400	0.000	90.440	11.550	13.087

Table 5-10: GLCM contrast statistics with large window from band 26, DIRSIG data

The results show that even with the increased window size, that data exhibit the same trend as the data in Table 5-5 and Table 5-6. The values are not the same as the previous evaluation, but comparatively, the mean value data still support the conclusion that

contrast between targets and background are appropriate in the DIRSIG images. The results of the evaluation in band 60 are shown in Table 5-11 and Table 5-12.

GLCM Contrast Statistics - AHI Data (large window, Band 60)				
	Min	Max	Mean	Standard Deviation
E700	0.000	145.095	35.708	26.023
E1400	0.000	123.400	36.033	26.243
N700	0.000	196.762	44.070	36.746
N1400	0.000	120.640	25.156	20.070

Table 5-11: GLCM contrast statistics with large window from band 60, AHI data

GLCM Contrast Statistics - DIRSIG Data (large window, Band 60)				
	Min	Max	Mean	Standard Deviation
E700	0.000	271.610	67.307	42.374
E1400	0.000	252.600	79.885	51.737
N700	0.000	135.333	25.607	20.350
N1400	0.000	70.600	9.817	8.628

Table 5-12: GLCM contrast statistics with large window from band 60, DIRSIG data

Again, these results show with the increased window size, that data exhibit the same trend as the previous GLCM evaluation in band 60. The contrast level is higher in the evening, DIRSIG data and lower in the midday data, compared to the AHI results. This supports the conclusions developed from the initial evaluation.

Spectral Co-occurrence Matrix analysis also provides a detailed look at the representation of contrast between disturbed soil and undisturbed soil. The difference between GLCM analysis and SCM analysis is that in SCM analysis the contrast is determined between two separate spectral bands, in an attempt to provide a measure of spectral contrast. The results from this metric are less easily interpreted. Extensive research has been accomplished in order to interpret GLCM analysis and declare that the contrast metric accurately predicts contrast in the image. SCM analysis is new to the field, being used at RIT only one other time on imagery that was spatially equivalent. While it is assumed that the SCM contrast metric calculates spectral contrast, thorough investigation to prove this, is not available in the literature. Pre-processing was accomplished in the same manner as the GLCM evaluations, the first evaluation using a smaller processing window, followed by an evaluation using a more “target-like” processing window. Two spectral bands were needed for the evaluation, so the band 26 subsets and band 60 subsets were used. Using a special adaptation of the GLCM program in ENVI (Scanlan, 2003), SCM processing was accomplished in the same manner as GLCM processing. At the outset of the evaluation, the expectation was that

the results should be no better or worse than the evaluations of bands 26 and 60 individually, as the bands are now be compared one to another. The results of the initial evaluation are shown in Table 5-13 and Table 5-14.

SCM Contrast Statistics - AHI Data				
	Min	Max	Mean	Standard Deviation
E700	0.000	764.444	153.278	132.435
E1400	0.000	604.333	141.301	130.185
N700	0.000	1135.778	174.061	180.325
N1400	0.000	1802.778	371.731	388.601

Table 5-13: SCM contrast statistics from AHI data

SCM Contrast Statistics - DIRSIG Data				
	Min	Max	Mean	Standard Deviation
E700	0.000	889.889	166.508	130.961
E1400	0.000	813.889	250.240	185.937
N700	0.000	1317.111	93.744	129.357
N1400	0.000	624.000	111.286	103.671

Table 5-14: SCM contrast statistics from DIRSIG data

These results match the expectations perfectly. As can be seen, these results exhibit the same trends as the GLCM evaluations in band 60. Contrast levels in the evening shots are higher in the DIRSIG renderings while contrast levels in the midday renderings are lower. These results do not point to any additional root causes for these discrepancies other than those that had been uncovered from the GLCM evaluations.

Secondly, this evaluation was performed using the enlarged processing windows, exactly the same as the secondary GLCM evaluations. These results are presented in Table 5-15 and Table 5-16.

SCM Contrast Statistics - AHI Data (large window)				
	Min	Max	Mean	Standard Deviation
E700	0.000	711.381	141.118	110.323
E1400	0.000	482.120	131.432	104.586
N700	0.000	690.619	160.571	136.736
N1400	0.000	1340.400	339.614	330.127

Table 5-15: SCM contrast statistics with large window from AHI data

SCM Contrast Statistics - DIRSIG Data (large window)				
	Min	Max	Mean	Standard Deviation
E700	0.000	597.381	155.158	114.626
E1400	0.000	626.040	228.853	164.219
N700	0.000	702.571	90.027	100.206
N1400	0.000	390.840	98.422	73.499

Table 5-16: SCM contrast statistics with large window from DIRSIG data

Again, the same trend is exhibited in these results as with the results from the initial SCM evaluation. In short, both SCM evaluations point to discrepancies in the data that were determined in the GLCM analysis. For future evaluations of this nature it is suggested

that the more useful measure is the GLCM analysis over a multitude of spectral bands. In addition, it may be useful to perform a wider analysis of background contrast by selecting a larger subset region and/or a much bigger processing window size. This evaluation attempted to tune the processing window to the size of the targets in the subsets, facilitating a reasonable target to background contrast evaluation. A better measure of overall background contrast might have been obtained by processing the data with a larger processing window.

5.6 R(x) Algorithm Performance

The final evaluation is a comparison between an anomaly detection algorithm's performance on the truth scene and the synthetic scene. Many "canned" anomaly detection algorithms exist today, but in an attempt to provide a more strenuous test, the R(x) anomaly detection algorithm was decided upon. This algorithm processes a multi or hyperspectral data set spectrally and spatially over a user defined processing window size. This is unique in that standard anomaly detection algorithms tend to process either spectrally or spatially, but not in a combined fashion. It is important to note that this algorithm is not designed specifically to address the mine detection issue; so excellent mine detection was not a requirement. The goal was to see if the algorithm would perform equally as well on both data sets, whether that be good or bad. Each truth and synthetic scene was pre-processed using a standard PC transform. Only the first five PC bands of data were input into the algorithm to eliminate very lengthy run times. Based on the results from the dimensionality analysis, this approach seems very reasonable. A pre-coded version of the algorithm written with an ENVI interface was used for all processing. The reader is encouraged to reference Barcomb (2004) for specific implementation of this algorithm in the ENVI environment. A generic processing window size of 21x21 pixels was used for all evaluations, with a target spatial shape defined as a 5x5 square set of pixels for the 700-foot data and a 3x3 square set of pixels for the 1400-foot data. The size of the processing window was chosen such that when the window was centered on a buried mine target the maximum amount of area was captured without including other targets from above or below. The target spatial shapes were

selected to tune the algorithm to the approximate spatial size of the buried mine areas in the truth imagery. In order to visually compare the results of the algorithm on the AHI and DIRSIG data accurately, the images have had a 98% image data threshold applied. Comparison images for all data sets are presented in Figure 5-27 through Figure 5-30.

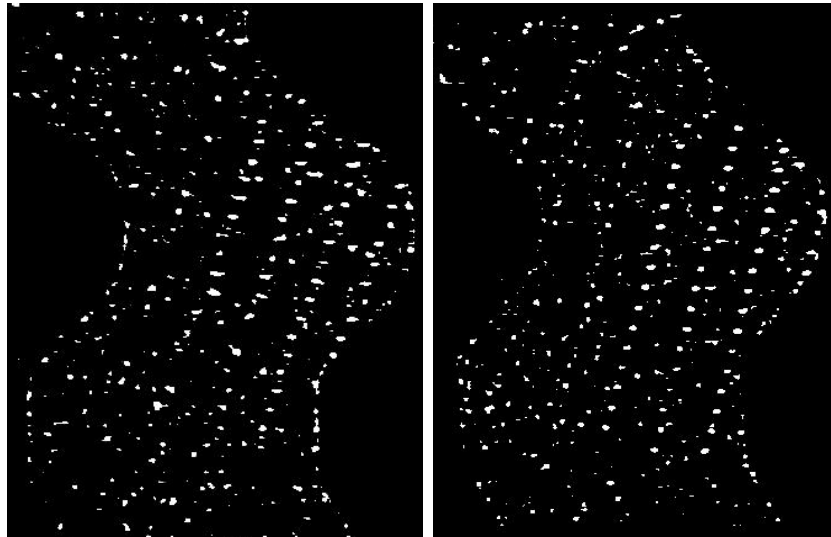


Figure 5-27: R(x) result images for (L) AHI data and (R) DIRSIG data, midday at 700 feet

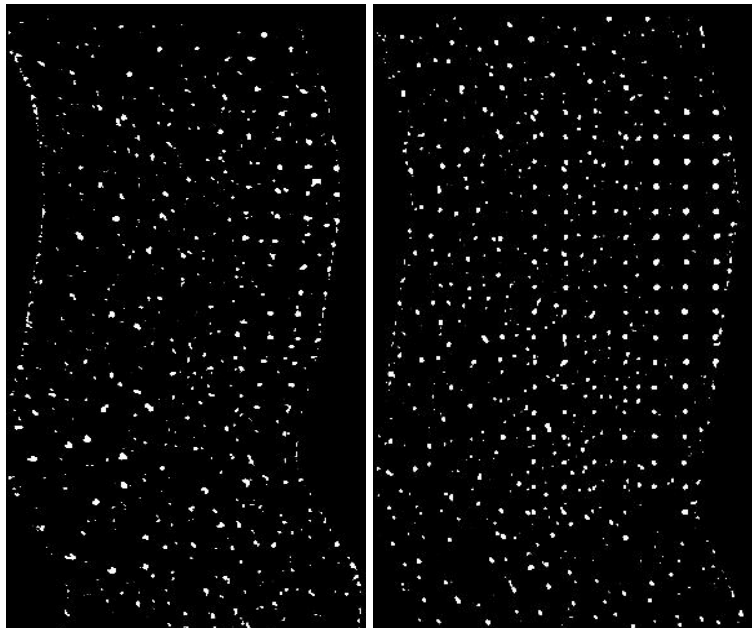


Figure 5-28: R(x) result images for (L) AHI data and (R) DIRSIG data, midday at 1400 feet

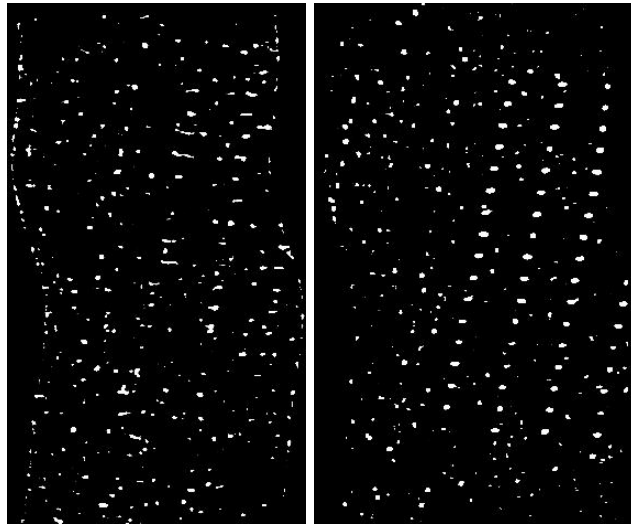


Figure 5-29: R(x) result images for (L) AHI data and (R) DIRSIG data, evening at 700 feet

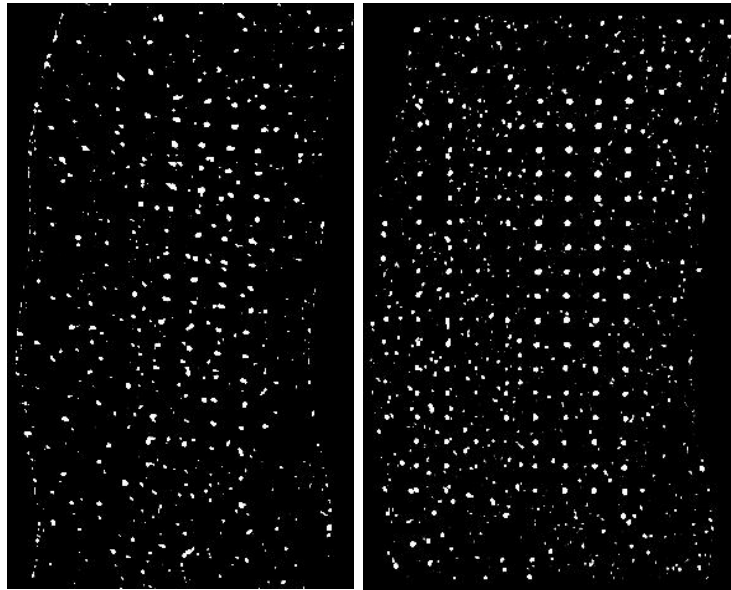


Figure 5-30: R(x) result images for (L) AHI data and (R) DIRSIG data, evening at 1400 feet

The result images are very encouraging. The numbers of false alarms are very significant in both the truth and synthetic data. While the results are not identical, they show that spectral and spatial clutter has been created that can adequately model the real world. The amount of clutter seen in the synthetic data, while appearing to be less than the truth, is of lower concern from a modeling standpoint because additional clutter objects can easily be inserted. The fact that these clutter objects approximate truth clutter quite well spatially and spectrally is the significant point. It was expected that mined areas would

be more noticeable in the synthetic data, purely due to the fact that it is modeled data. Some of this can be explained. All of the buried mine areas were created to have a very regular circular shape, where true buried mine areas are quite irregularly shaped. Adding more spatial irregularity to these areas would have a dramatic effect in the algorithm's ability to detect. This is precisely the reason behind the noticeable discontinuity between the results in the evening, 700-foot data. As mentioned before, the target's spatial structure was input to the algorithm as a regular square of pixels, which will more closely match the synthetic data than real world. From a scene building prospective, this scene more closely represents the truth to a target detection algorithm than any other thermal scene built at the DIRS lab to date. That being said, this scene is only the first step to exact statistical representation of the truth and will be improved upon.

In an attempt to add a quantitative side to this evaluation, Receiver Operator Characteristic (ROC) curves were generated for each data set. Generating the curve set was more difficult than originally anticipated. Specifically, it was extremely difficult to select known target pixels in the truth imagery. Selection of target pixels was done purely by eye, which leads to an unavoidable source of error. Also, the image sizes used for the evaluation were different, leading to different numbers of pixels being used to determine final probabilities. Subsections of the imagery were used in an attempt to mitigate this effect, however the problem could not be completely resolved. Another source of unavoidable error comes from the inexact reproduction of the AHI flight lines. The flight lines were intended to be similarly matched, but not reproduced, leading to slightly different geometric presentations of target areas and the entire scene in general. Given these limitations, the ROC curves are presented in Figure 5-31 through Figure 5-34 as a baseline to be used for comparison to future versions of this scene.

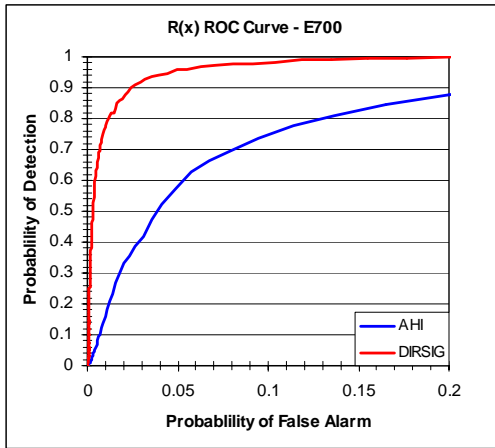


Figure 5-31: R(x) ROC curve, evening at 700 feet

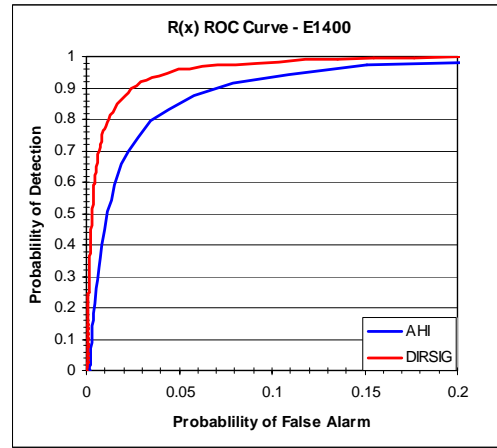


Figure 5-32: R(x) ROC curve, evening at 1400 feet

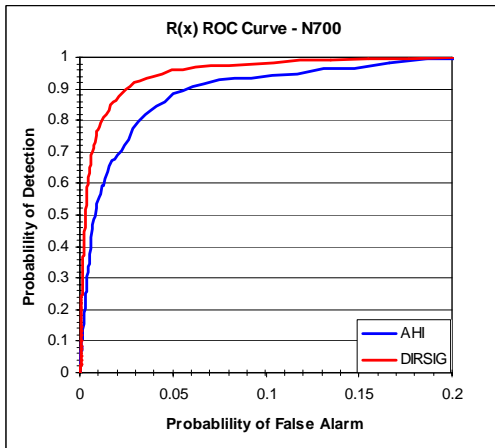


Figure 5-33: R(x) ROC curve, midday at 700 feet

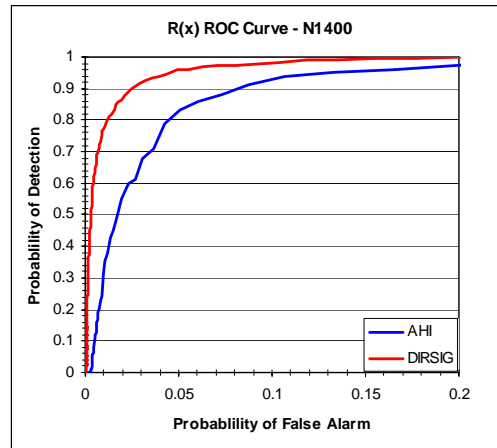


Figure 5-34: R(x) ROC curve, midday at 1400 feet

As pointed out previously from the imaged-based R(x) results, the evening, 700-foot data shows the largest discrepancy between the synthetic and truth's evaluation results. Regardless of the difficulties in generating these curves, the relative difference between the evening, 700-foot data and the others is the important point. As discussed previously, the dissimilarity in the data stems predominantly from the spatial shape of the target regions. These results suggest that creating irregularly shaped target regions should be a considered a first step for the next cyclical revision of this scene. However, these results also suggest that clutter is being adequately represented throughout all versions of the scene and that further detailed evaluation in this area is unnecessary.

6.0 Conclusions and Recommendations

Four unique renderings of a synthetic surface and buried landmine scene have been created and compared to calibrated truth data. This work has demonstrated that DIRSIG is capable of accurately modeling a representative LWIR scene complete with spectrally and spatially varying clutter sources, even when lacking fundamental modeling parameters. New techniques have been developed and used to assist the scene builder in deriving accurate input data for the model, such as synthetic correlated noise generation and emissivity curve derivation. Moreover, this work has shown the need for further study in the area of temperature variability modeling combined with emissivity texturing. In addition, a description of the Microscene experiment and the coordination needed to undertake a collection of that magnitude has been presented. Hopefully this experiment can be used as a model for future collects.

6.1 Cyclical Modeling

Scene modeling is cyclical in nature; a first version is created and evaluated, and then upgrades are determined and implemented. The new scene is rendered and the process repeats. While exact replication of the truth scene may never be perfectly attainable, striving to understand the error takes us a few steps closer to equality. This project has gone through two full cycles of revision after the initial foundation was laid. With each revision, the focus for scene improvement changed, beginning with underlying geometry and problematic DIRSIG concerns, ending up at targeted phenomenology issues. This revision process is absolutely essential for the creation of any valid scene since it is impossible to fully investigate intricate scientific concerns if confidence in running the model has not been previously attained in earlier revision cycles. Figure 6-1 attempts to show the cyclic process for this project from a visual prospective. The process depicted is the revision cycle for the evening, 700-foot rendering, focusing on spectral band 25.

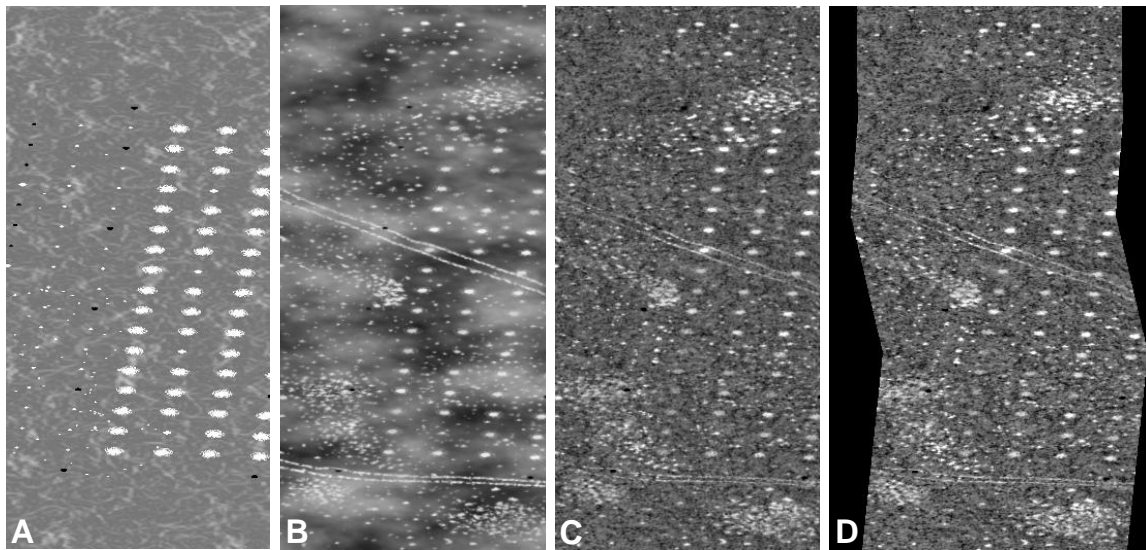


Figure 6-1: Pictorial description of the cyclical scene building process used for the landmine scene

Part A is a rendering after the initial foundation of the scene had been laid. This version included first cut geometry and placement for all scene targets, the first attempt at including a bump map and very basic texturing for undisturbed soil. Evaluation of this version pointed out that there were some sizing issues with certain targets, a need to include much better texture, and something that can't been seen in the image, a need to address temperature variability and improve material parameters. What was good about the first rendering was the geometric placement of all the targets in the scene as well as the initial aircraft flight lines and sensor modeling. The part B rendering addressed the concerns from the part A rendering. Sizing issues were eliminated, better texture was developed, realistic clutter was included and some of the temperature distribution concerns were dealt with. From a visual standpoint, one of the most dramatic improvements that enhance the realism of the rendering is the inclusion of inter-pixel blur. In fact, the project dove deeper into aspects of sensor modeling than originally anticipated, but was without question the right way to go. Version B was the first version delivered to the Army MURI team in December 2003. After the first complete revision, the focus areas for improvement began to target more phenomenological concerns. Specifically, work was needed in perfecting the target/clutter/background contrast, refining atmospheric constituents, and investigating ways to improve the combined temperature distribution and emissivity distribution effect. The results of these efforts are

apparent in version C. The issue of emissivity texture was greatly improved, contrast between targets and background were updated and atmospheric modeling issues were addressed. Version D is what I consider to be half a revision. This version is simply version C with the added aircraft roll phenomenology. After versions C and D were rendered, the point had been reached where only detailed metric analysis could provide concrete direction for further development. Versions C and D were delivered to the Army MURI team in April 2004. The scene in its current form is now ready for the next full revision cycle, given the results of the analysis presented here. Readily apparent from these results is the need for research pertaining to LWIR texturing and spatial temperature distribution.

6.2 Scene Improvements

The objective of this research was to develop the landmine scene to the point in which detailed analysis was required to determine the revisions that should be made for the next cyclical iteration. Based on analysis of the current scene, temperature profiles for buried mine areas need to be obtained and utilized. The current material map incorporates two different types of soil representing the buried mine target as a whole. These include the soil directly above the buried mine and the disturbed soil surrounding the buried mine. If temperature profiles for these types of soils were obtained, the avenues for incorporation into the scene are available. Secondly, the locations of the buried mine areas in the scene should be varied, reducing some of the geometric regularity of the minefield. In addition, the spatial shape of each individual buried mine area should be changed such that they appear more random. By altering the basic form of these regions throughout the scene, a more natural and realistic look will be obtained. Third, additional efforts should be made to obtain target and background spectra that encompass the full spectrum, from visible to long wave infrared. Not only will the addition of these spectra allow the scene to be rendered in other regions of the spectrum, but also the visible regions of the emissivity curves can be used to drive the solar absorption for each material and allow the solar absorption to be spatially located based on the scene's emissivity map. This would be invaluable for beginning to resolve the

issues with the spatial distribution of temperature data. If full-spectrum emissivity curves cannot be obtained, obtaining a broadband, overhead visible image of the scene could also be used to generate a solar absorption map, which could spatially locate this material property throughout the scene, also solving some of the issues noted with the spatial distribution of temperatures. Lastly, better sensor noise data should be obtained and incorporated. While the noise derivation process has done an excellent job of allowing the inclusion of spectrally correlated noise that approximates true sensor noise, real measures of sensor noise, such as a full-field blackbody scan, would allow not only spectrally correlated noise, but spectrally and spatially correlated noise. This data would allow the addition of detector striping, as observed in the visual image comparison.

6.3 Fundamental LWIR Modeling Improvements

From an overarching modeling point of view, there are a few areas of research that would greatly enhance DIRSIG's ability to generate statistically and visually accurate scene in the LWIR. First, a method for incorporating a true sensor point spread function into the rendering process should be developed. This addition to the DIRSIG program would eliminate additional image processing steps needed to fully incorporate a sensor's point spread function. Secondly, research should continue in developing methods to improve thermal variability. As mentioned previously, using broadband visible data to drive a solar absorption map would greatly enhance thermal variability and its spatial distribution, since solar absorption is a key parameter in the development of a material's diurnal temperature curve. Unfortunately, the truth data set used in this work did not have any accessible overhead visible imagery. However, when determining ground truth data requirements for future modeling, it would be prudent to ensure that broadband visible and broadband thermal overhead imagery of the scene in question are collected. Another area of investigation to improve temperature variability scene-wide would be to incorporate a range of values for material parameters such as thermal conductivity, specific heat, or mass density, rather than a single value. Research would need to be performed to determine a realistic range for each parameter, however it seems more realistic to expect a range of values across a landscape. Once this range has been

determined for each of the varying parameters, the modeler could then determine the number of variations for each parameter. A program could run the thermal model off-line for each material combination producing a multitude of diurnal temperature curves for soil in the landscape. Once all the diurnal curves are created, spatial distribution throughout the scene could be accomplished similarly to the texturing process. A gray level image would drive the placement of each temperature curve, using the Z-score method or any other texturing method. The difficulty would be to obtain a gray level image that represents the true distribution of these soil parameters in the truth scene. Regardless of the method developed to improve temperature variability, great care must be taken to ensure that temperature data is correctly reconciled with emissivity data. It must be ensured that a correctly modeled temperature is not being masked by an incorrectly placed emissivity curve. An avenue for ensuring this reconciliation exists by using images of the truth scene taken at multiple times of day to observe the time when a thermal crossover period occurs. At one of these periods, thermal variation would be at a minimum and spatial emissivity distribution over the landscape can be observed. If the data are overhead images encompassing the scene of interest, these thermal crossover images can be used to drive the modeled scene's emissivity map, and allowing reconciliation between temperature distribution and emissivity distribution to occur. Research into these areas is key in furthering LWIR scene modeling.

6.4 Summary

This project has demonstrated that DIRSIG is fully capable of producing a statistically accurate LWIR scene complete with spectrally and spatially varying clutter sources, even when lacking fundamental modeling parameters. This project was oriented towards the development of target and anomaly detection algorithms. With this in mind, future work should be to fully investigate methods that will improve LWIR scene modeling from a statistical point of view. A few of these potential research areas are presented in this thesis. These additional steps will provide a more robust training ground for algorithms. Hopefully, this work can provide a solid foundation for statistically accurate LWIR scene creation.

7.0 References

- AHI Webpage, “AHI – Airborne Hyperspectral Imager for Remote Sensing in the longwave 8 – 12 micron (thermal) Infrared”, accessed online at <http://www.higp.hawaii.edu/ahi/>, (Jul 2003).
- Ark Thermal Insulation and Engineering Trading Company, “Non-Metals Thermal Properties”, accessed online at <http://www.arkthermal.com/non-metals.doc>, (Dec 2003).
- Army, Department of the, “Mine/Countermining Operations”, *Field Manual 20-32*, (2002).
- Baraldi, A. and Parmiggiani, F., “An Investigation of the Textural Characteristics Associated with Gray Level Cooccurrence Matrix Statistical Parameters,” in IEEE Transactions on Geoscience and Remote Sensing, Vol. 33, No. 2, pp. 293–304, (Mar 1995).
- Barcomb, K., “High-resolution, Slant-angle Scene Generation and Validation of Concealed Targets in DIRSIG”, M.S. Thesis, Rochester Institute of Technology, (2004).
- Berk, A., et al., “MODTRAN4 User’s Manual”, Air Force Research Laboratory, Space Vehicles Directorate, AFMC, Hanscom AFB, (1999).
- Bonsor, K., “How Landmines Work,” accessed online at *How Stuff Works.com*, <http://www.howstuffworks.com/landmine.htm>, (2001).
- Boras, I., Malinovec, M., Stepanic Jr., J., Svaic, S., “Detection of underground objects using thermography,” Proc. of the 15th World Conference on Non-Destructive Testing, (Oct 2000), accessed online at <http://www.ndt.net/article/wcndt00/papers/idn106/idn106.htm>, (Sep 2003).
- Brown, S., “DIRSIG User’s Manual, Release 3.5”, publication of the Digital Imaging and Remote Sensing Laboratory, (Mar 2003).
- Brown, S., Raqueño, R., Schott, J., “Incorporation of bi-directional characteristics into the Digital Imaging and Remote Sensing Image Generation model”, in Proc. of the Ground Target Modeling and Validation Conference, (Aug 1998).
- Brown, S. and Schott J., “Characterization techniques for incorporating backgrounds into DIRSIG”, in Proc. SPIE Vol. 4029, pp. 205–216, (2000).

- Cremer, F., de Jong, W., Schutte, K., “Infrared polarization measurements and modeling applied to surface-laid antipersonnel landmines”, in Optical Engineering Vol. 41, No. 5, pp. 1021–1032, (2002).
- DCS Corporation, “AIRSIM thermal signature prediction and analysis tool model Assumptions and analytical foundations”, DCS technical note 9090-002-001, (1991).
- DePersia, A., Bowman, A., Lucey, P., Winter, E., “Phenomenology considerations for hyperspectral mine detection”, in Proc. SPIE vol. 2496, pp. 159–167, (1995).
- DePersia, A. et al., “ARPA’s hyperspectral mine detection program”, in Proceedings of the Third International Symposium on Spectral Science Research (ISSSR), (1995).
- DIRSIG Homepage, accessed online at <http://www.cis.rit.edu/~dirsig>, (Jul 2003).
- ENVI User’s Guide, ENVI Version 3.2, Better Solutions Consulting Limited Liability Company, (July 1999).
- Haralick, R., Shanmugam, K., Dinstein, I., “Textural Features for Image Classification”, in IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, No. 6, pp. 610– 621, (Nov 1973).
- Hong, S., Miller T., Borchers, B., Hendrickx, J., Lensen, H., Schwering, P., van den Broek, S., “Land mine detection in bare soils using thermal infrared sensors”, in Proc. SPIE vol. 4742, pp. 43–50, (2002).
- Incropera, F.P. and DeWitt D.P., *Fundamentals of Heat Transfer*, Wiley and Sons (New York), (1981).
- Janssen Y., de Jong, A., Winkel, H., van Putten, F., “Detection of surface laid and buried Mines with IR and CCD cameras, an evaluation based on measurements”, in Proc. SPIE vol. 2765, pp. 448– 459, (1996).
- Joseph, D.J., “DIRSIG: A Broadband Validation & Evaluation of Potential for Infrared Imaging Spectroscopy”, M.S. Thesis, Rochester Institute of Technology, (1998).
- Kahle, A.B. and Alley, R.E., “Separation of Temperature and Emittance in Remotely Sensed Radiance Measurements”, in Remote Sensing of the Environment, vol. 42, pp. 107–111, (1992).

- Kealy, P.S. and Hook, S.J., “Separating Temperature and Emissivity in Thermal Infrared Multispectral Scanner Data: Implications for Recovering Land Surface Temperatures”, in IEEE Transactions on Geoscience and Remote Sensing, vol. 31, no. 6, pp. 1155 – 1164, (Nov 1993).
- Khanafer, K. and Vafai, K., “Thermal analysis of buried land mines over a diurnal cycle”, in IEEE Transactions on Geoscience and Remote Sensing, vol. 40, no. 2, pp. 461–473, (2002).
- Kiang, R.K., “Utilizing spatial features in classifying high-resolution imagery data”, in Proc. SPIE vol. 4725, pp. 267– 274, (2002).
- Korb, A., Dybwad, P., Wadsworth, W., and Salisbury, J., “Portable Fourier transform Infrared spectroradiometer for field measurements of radiance and emissivity,” in Applied Optics, Vol. 35, No. 10, pp. 1679–1692, (Apr 1996).
- Kraska, T.A., “DIRSIG, (Digital Imaging and Remote Sensing Image Generation) Model: Infrared Airborne Validation and Input Parameter Analysis”, M.S. Thesis, Rochester Institute of Technology, (1996).
- Larson, R.E. and Edwards, B.H., *Elementary Linear Algebra, Third Edition*, D.C. Heath and Company, (1996).
- Liao, W.J., Chen, D.H., Baertlein, B.A., “Detection of Land Mines in Multi-Spectral and Multi-Temporal IR Imagery”, accessed online at <http://www.ee.duke.edu/~lcarin/DeminingMURI/DeminingMURIpublications.html>, (Jul 2003).
- Lucey, P.G., and Winter, E.M., “Requirements for calibration of focal plane arrays for imaging spectrometers”, in Proc. SPIE vol. 3498, pp. 305–310, (1998).
- Maksymonko, G., Ware, B., Poole, D., “A characterization of diurnal and environmental effects on mines and the factors influencing the performance of mine detection ATR algorithms”, in Proc. SPIE vol. 2496, pp. 140– 151, (1995).
- Mason, J., Schott, J., Salvaggio, C., Sirianni, J., “Validation of contrast and phenomenology in the Digital Imaging and Remote Sensing (DIRS) lab’s Image Generation (DIRSIG) model”, in Proc. SPIE vol. 2269, pp. 622–633, (1994).
- McGovern, M. and Aponte, H., “Analysis of IR signatures and buried anti-tank Landmines”, in Proc. SPIE vol. 4394, pp. 237–245, (2001).
- McKeown, D., “Wildfire Airborne Sensor Program (WASP) Project Overview”, presentation given at the Rochester Institute of Technology’s Industrial Associates Meeting, (May 2003).

- MISI Homepage, *Modular Imaging Spectrometer Instrument*, accessed online at <http://www.cis.rit.edu/research/dirs/research/misi.html>, (July 2003).
- Nivelle, F. and Lhomme P., “Detection of land mines with passive IR and mmW imaging Sensors”, in Proc. SPIE vol. 3079, pp. 614– 624, (1997).
- Norwegian People’s Aid, *Landmine Database*, accessed online at http://www.angola.npaid.org/minelist_all_database.htm, (2003).
- Peebles Jr., P.Z., *Probability, Random Variables and Random Signal Principles (Fourth Edition)*, The McGraw-Hill Companies Inc., (2002).
- Pitts, D.R. and Sissom, L.E., *Schaum’s Outline of Theory and Problems of Heat Transfer (Second Edition)*, The McGraw-Hill Companies Inc., (1998).
- Rankin, D.K., “Validation of DIRSIG an Infrared Synthetic Scene Generation Model”, M.S. Thesis, Rochester Institute of Technology, (1992).
- Reed, I.S., and Yu, X., “Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution”, in IEEE Transactions on Acoustics, Speech and Signal Processing, vol. 38, No. 10, pp. 1760–1770, (Oct 1990).
- Russell, K., McFee J., Sirovyak, W., “Remote performance prediction for infrared imaging of buried mines”, in Proc. SPIE vol. 3079, pp. 762–769, (1997).
- Salacain, J.M., “Application of Imaging System Geometric Models to a Synthetic Image Generation System”, M.S. Thesis, Rochester Institute of Technology, (1995).
- Scanlan, N., “Comparative Performance Analysis of Texture Characterization Models in DIRSIG”, M.S. Thesis, Rochester Institute of Technology, (2003).
- Schott, J.R., *Remote Sensing: The Image Chain Approach*, Oxford University Press, NY, (1997).
- Schott, J.R., S.D. Brown, and M.J. Richardson, “The Role of Physical Modeling and Scene Simulation in Support of Space Based Remote Sensing”, presented at the Core Technologies for Space Systems Conference, Colorado Springs, CO, (Nov, 2001).
- Schott, J.R., Brown, S.D., Raqueño, R.V., Gross, H.N., and Robinson, G., “An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development”, Canadian Journal of Remote Sensing, Vol. 25, No. 2, pp. 99–111, (June 1999).

- Schott, J.R., “Advanced Hyperspectral Algorithms”, class notes, Rochester Institute of Technology, Rochester, New York, (Spring quarter, 2003)
- Simard, J.R., “Improved landmine detection capability (ILDC): Systematic approach to the detection of buried mines using passive IR imaging”, in Proc. SPIE vol. 2765, pp. 489–500, (1996).
- SIMS Brochure, “Simulant Mines (SIMS)”, distributed by the Department of the Army Mines, Countermine, and Demolition Program Office, (2003).
- SIMS Homepage, “Simulant Mines”, accessed online at http://ccsweb.pica.army.mil/2cmine/2_cmine_.htm, (Jul 2003)
- Smith, A., Kenton, A., Horvath R., Nooden, L., Michael, J., Wright, J., Mars, J., Crowley, J., Sviland, M., Causey, S., Lee, D., Williams, M., Montavon, K., “Hyperspectral mine detection phenomenology program”, in Proc. SPIE vol. 3710, pp. 819–829, (1999).
- Smith, H.P.J., et al., “FASCODE – Fast Atmospheric Signature Code (Spectral Transmittance and Radiance)”, AFGL-TR-78-0081, Air Force Geophysics Laboratory, Hanscom AFB, (1978).
- Stein, D., Schoonmaker, J., Coolbaugh, E., “Hyperspectral Imaging for Intelligence, Surveillance, and Reconnaissance”, Space and Naval Systems Warfare Center (SSC) San Diego Biennial Review 2001, SSC San Diego TD 3117, pp. 108–116 (Aug 2001).
- Swanson, N., Welker, H., Blume, B., Rish, J., “Application of thermal imaging to remote airfield assessment”, in Proc. SPIE vol. 3079, pp. 819–830, (1997).
- White, R.A., “Validation of RIT’s DIRSIG Model – Reflective Region”, M.S. Thesis, Rochester Institute of Technology, (1996).
- Winter, E. et al., “Experiments to support the development of techniques for hyperspectral mine detection”, in Proc. SPIE vol. 2759, pp. 139–148, (1996).
- Winter, E., “AHI at Arid Army Test Site”, presentation given as part of arid test site data collection effort, (Jul 2003).
- Yu, X., Reed, I.S., Stocker, A.D., “Comparative Performance Analysis of Adaptive Multispectral Detectors”, in IEEE Transactions on Signal Processing, Vol. 41, No. 8, pp. 2639–2686, (Aug 1993).